

What is it to understand?

Andrei Ionuț Mărășoiu
București, România

B.A. Theoretical Philosophy and Logic, University of Bucharest 2008
M.A. History and Philosophy of Science, University of Bucharest 2010
M.A. Neurophilosophy, Georgia State University 2013

A Dissertation presented to the Graduate Faculty of the University of Virginia in Candidacy for the
Degree of Doctor of Philosophy

Department of Philosophy

University of Virginia
April 2019

Dissertation advisor:

Brie Gertler

Dissertation committee members:

James Cargile

Paul Humphreys

Sarah Corse

Table of contents

Dissertation abstract.....	6
Acknowledgments.....	8
Dissertation introduction.....	9
1. Three approaches.....	9
2. Everyday uses of the word “understanding”.....	11
3. Attribution.....	14
4. Variety.....	18
5. Decision and reality in ascertaining understanding.....	20
6. The epistemology of understanding.....	23
7. The phenomenology of understanding.....	26
8. How the epistemology and the phenomenology of understanding are related.....	28
Chapter I Intellectual virtues and biased understanding.....	31
1. Understanding.....	32
2. Virtues and skills.....	33
3. Biases.....	35
4. Carter and Pritchard’s challenge.....	37
5. Reflective endorsement.....	42
6. Acts of virtue and skill.....	45

7. An example.....	49
8. Conclusion.....	52
Chapter II Mathematical understanding and “What if things had been different?” questions.....	53
1. Introduction.....	53
2. Understanding as knowledge of dependencies.....	54
3. Two simple cases of mathematical understanding.....	58
4. A modal approach to answering “What if things had been different?” questions.....	61
5. Grounding to the rescue?.....	63
6. Different models.....	67
7. Conclusion.....	71
Chapter III Experience and understanding.....	72
1. Introduction.....	72
2. Setting the stage.....	73
3. Grasping by experiencing.....	76
4. Is unconscious grasp as much as possible?.....	79
5. Inference, grasp, and conscious experience.....	81
6. Insight and understanding.....	85
7. Intuition and insight.....	88
8. Creativity and understanding.....	91
9. Modes of understanding.....	94

10. Conclusion.....	101
Chapter IV Understanding: conscious experience and epistemic norms.....	103
1. Introduction.....	103
2. Bengson on understanding.....	107
3. Wittgenstein's question.....	110
4. Intuitions and the epistemic norms of understanding.....	114
5. Conscious experiences associated with understanding.....	118
6. Understanding, explanation, models and skills.....	123
7. We understand by attending.....	126
8. How are phenomenal and epistemic aspects of understanding related?.....	129
9. Defending the premises.....	131
10. Addressing Wittgenstein's question.....	135
11. Conclusion.....	138
General conclusion.....	140
1. Summary.....	140
2. The unity of understanding.....	142
3. Consequences of the dissertation for several philosophical debates.....	144
4. Open questions.....	149
References.....	153

Dissertation abstract

How are epistemic and phenomenal properties of understanding related? In my dissertation, I offer an answer to this question, relying on an account of how understanding gains its positive epistemic status and on an account of what conscious experiences are associated with understanding. Chapters I and II discuss epistemic aspects of understanding. Chapter III concerns experiences associated with understanding, and classifies them according to the modes in which understanding occurs to us. Then, in Chapter IV, I address the epistemic bearing of these types of experiences for understanding.

In Chapter I, I defend the virtue epistemology of understanding – whose most well-known proponent is Zagzebski (2001) –, against a skeptical objection which Carter and Pritchard (2016) raise considering the ubiquity of cognitive biases. In Chapter II, I argue that two instances of mathematical understanding are counterexamples to Grimm's (2014) view of understanding, according to which we only understand what we can explain, we can only explain a phenomenon if we grasp what its occurrence depends on, and we identify such dependencies by answering “What if things had been different?” questions. (I remain neutral about which of Grimm's claims should be abandoned.) Jointly, Chapters I and II support a limited pluralism about what epistemic norms are met by states of understanding.

In Chapter III, I build on the work of Bourget (2017) and Lynch (2017) to argue that understanding occurs in one of three modes. Each mode has a kind of conscious experience associated with it. Moments of coming to understand are conscious in insights. Discursively exploiting an understanding one already possesses is conscious in intellectual intuitions. And non-discursively exploiting an understanding one already possesses is manifest in experiences of fluency.

Previous chapters lead up to the question Chapter IV addresses: how do conscious experiences

meet the epistemic norms of understanding? The most promising attempt to address this question, I argue, identifies insights, intuitions, and experiences of fluency as kinds of experience associated with understanding. And it specifies that the use of models to represent phenomena, the ability to explain those target phenomena, and the manifestation of our cognitive skills, are possible epistemic norms for understanding. The question of how phenomenal and epistemic aspects of understanding are related may then be replaced with a more precise one, namely, how these specific experiences are related to meeting these specific epistemic norms.

In Chapter IV, I hypothesize that conscious experiences are related to meeting the norms of understanding is *via* attention: we understand by attending. In insights, we shift the pattern of what we attend to, gaining the ability to explain what we sought to understand. In intuitions, how we attend to phenomena understood is framed by models representing those phenomena. And experiences of fluency, including ease and familiarity in how we approach phenomena we tacitly understand, are associated with the manifestation of cognitive skills on that particular occasion.

However, in Chapter IV I also argue that the correlations between insights, intuitions, and experiences of fluency, and, respectively, meeting the epistemic norms of understanding, can only be supported with claims that have at best the modal strength of natural necessity, and are metaphysically contingent. That is, we have insufficient evidence to claim that there are any conscious experiences that the occurrence of understanding is partly constituted by. And we have insufficient evidence to claim, with the strength of metaphysical necessity, that understanding occurs in virtue of our experiences, or is ultimately causally explained by our having such experiences.

Acknowledgments

This dissertation is the fruit of wonderful questions, comments and suggestions I have received over the past few years. I am especially grateful to Brie Gertler, Jim Cargile and Paul Humphreys. I owe a debt of gratitude to my previous advisors and mentors as well: Gabriel Săndoiu, Mircea Dumitru, Mircea Flonta, and George Graham.

It cannot be overemphasized how fortunate I have been to receive input from other friends and colleagues too. They include, in alphabetical order of their last name: Sorin Bangu, Tal Brewer, Lauren Callahan, Jeff Carroll, Sarah Corse, Elena Costea, Jim Darcy, Matt Duncan, Sandy Dwyer, Bogdan Florea, Michele Friend, Torrance Fung, Geoff Gordon, Nora Grigore, Jaakko Hintikka, Zac Irving, Petrișor Ivan, Hans Kamp, Derek Lam, Larry Laudan, Guy Longworth, Eddy Nahmias, Laurie Paul, Michael Pendlebury, Mark Sainsbury, Gabriel Sandu, Jim Seitz, Paul Teller, Silviu Velica, Dan Weiskopf, Erin Westgate, and many others. Many thanks are also due to anonymous journal reviewers who have considered parts of this manuscript or my related ideas.

Several drafts of chapters were given as talks or circulated as journal manuscripts. I am grateful to audiences at the Joint Session of the Mind Association and Aristotelian Society in 2017 and 2018, the North Carolina Philosophical Society meeting of 2016, the CLMPS 2015 and HOPOS 2016 meetings, many PAD workshops at UVa, and a seminar series convened at the Institute for Research in Humanities in Bucharest over the summer of 2017.

Financial support was provided by the Jefferson Scholars Foundation through a John S. Lillard fellowship, and the University of Virginia through a Dissertation Completion Fellowship. Last but not least, this work wouldn't have been possible without the support and encouragement of my family.

Dissertation introduction

In this Introduction, I draw the contours of the intellectual kind of understanding that I will investigate in the dissertation. I will seek to do justice to three approaches, according to which, respectively, conscious experiences associated with understanding are quite varied, the epistemic norms governing understanding may also vary, and yet there might be metaphysical unity to understanding notwithstanding. I show that these three approaches to understanding need not be mutually inconsistent, and articulate the question of how phenomenal and epistemic aspects of understanding are related to each other.

1. Three approaches

Consider the question “What is understanding?” Moravcsik thinks this is a purely metaphysical question, inquiring into the nature of a specific kind of state of mind:

The constitutive question; i.e., what is understanding, is distinct from the epistemological issue of how we know when we ourselves or someone else has achieved understanding. (Moravcsik 1979, p. 209)

In contrast, Cooper thinks philosophers need to make epistemological decisions as to which epistemic norms a state of mind should satisfy in order to count as understanding:

It is the task of the normative epistemology of understanding to adjudicate between competing conceptions of understanding. (Cooper 1994, p. 15)

And Goldfarb voices skepticism: the variety and irregularity of how we use the word “understanding” and its cognates prevents any fruitful generalization.

Wittgenstein asks us to look in detail at the range of our practices relevant to an ascription of understanding. We find an enormous variety of considerations that can enter, a dependence on context that is impossible to describe accurately by any general rules, a lack of uniformity in mental accompaniments. In individual cases we have stories to tell to justify our ascriptions, but there is no uniform feature that pulls the various cases together. (Goldfarb 1992, pp. 112-113)

These three approaches to understanding have received little attention recently. However, addressing them seems needed for any philosophical account of understanding to even get off the ground. The three approaches vary widely, and one may well think they are jointly inconsistent. For variety in what we *term* understanding may resist any unified account. And even if it didn't, perhaps that unified account would be merely *stipulative* – a conceptual construction – rather than answering to the reality of our life of thought.

In this Introduction, I draw the contours of the intellectual kind of understanding that I will investigate in the dissertation. I will seek to do justice to what is right in these three approaches, and to show why they are not mutually inconsistent when properly construed. Sections 2-5 are devoted to this task. In Sections 6-8, I summarize the dissertation to follow, focusing on the central question I pursue, namely, how phenomenal and epistemic aspects of understanding are related.

2. Everyday uses of the word “understanding”

How are the word “understanding” and its cognates used? Here are a few examples:

Hans understands the limits of Montague grammar all too well.

I understand that you're going to the countryside in June.

I understand what you're going through.

Marcia never understood why Caesar was stabbed.

Teddy understands when the time is right to leave the party.

Maya is such an understanding person.

Do you understand what I'm saying? Leave the dog alone.

We couldn't understand each other had we not lived through the same ordeals.

Jack lacks any understanding of how to teach algebra to high schoolers.

These uses are quite heterogeneous.¹ However, we can classify them into at least five different types of uses: perceptual, practical, emotional, linguistic and intellectual understanding. Let me illustrate each:

In order to *perceptually* understand that the house is on fire, we have to segment the visual scene around us relative to concepts we possess.

The *practical* understanding the ballerina has of her performance, or the cyclist of how he can win the race, are matters of know-how that may – but, crucially, need not – require beliefs, let alone

¹ This heterogeneity has led Goldfarb (1992) to think there is no useful generalization to be had about understanding. I return to the issue in Section 4 below.

knowledge or reflective insight.

Emotional understanding is different too. Many of our friends and acquaintances are said to be understanding on account of the empathy they may manifest to our plight. Part of a proper appreciation of religious life may require that we educate ourselves to have the right feelings of care and self-effacement toward our neighbor.

Linguistic understanding requires none of this. We may understand what someone says simply when we know what they mean to be saying. Further conditions may be added, such as awareness of the circumstances that made someone say what they did, or familiarity with one's mannerisms and interpreting their utterance in light of them, or considering the various *minutiae* that may make one's intervention seem tone-deaf, witty, or in self-conscious disavowal of the rules of etiquette.

In what follows, I leave these varieties aside. Rather, I will be concerned with clarifying *intellectual* understanding.² Of the examples given above, only two straightforwardly qualify as intellectual understanding:

Hans understands the limits of Montague grammar all too well.

² Here is an example mixing the types of use listed above: understanding the role played by religious motifs in Chagall's early work. There is a *practical* aspect: looking over those paintings, knowing how to identify the religious elements, and succeeding in identifying their significance. In appreciating the painting, one also needs to be *emotionally* attuned to the viewing experience – to feel as one should. And there is an *explanatory* element: in understanding the role religious motifs play in Chagall's early work, one should be able to (at least in part) explain that role to others who may lack one's particular perspective on Chagall's work. These elements are intertwined in the inquiry one initiates. The example matters because *many* cases of what we call “understanding” may be like this: mixing distinct types of uses in a way whose specificity can only be appreciated in a particular context. In deciding to explore intellectual understanding, I introduce an idealization of everyday word usage. I do so in order to relate everyday uses to what the various philosophical literatures about understanding target.

Marcia never understood why Caesar was stabbed.

We could add more, e.g.:

It's hard to understand why Althusser murdered his wife.

When will we finally understand the effects of socialism within the Iron Curtain?

The perceptual, practical, emotional, linguistic and intellectual uses of the word “understanding” seem so different that I doubt they depict the same state of mind. Focusing on distinct kinds of uses of the word “understanding” has a better chance to depict recognizable states of mind. Relying on the distribution of different uses of the word “understanding,” Franklin (1981) argues that there is a cluster of uses (roughly corresponding to what I call “intellectual understanding”) that depict a state of mind with positive epistemic status, similar to but also distinct from knowledge and belief.³ This is one reason why, in what follows, I will only be concerned with intellectual understanding.

³ As Baumberger et al. (2016, pp. 29-30) note, we can use the verb “to understand” adjoined to different grammatical categories: understanding a person or language, understanding *that* Napoleon lost at Waterloo, understanding *how* a steam engine works, understanding *why* floods may follow the first warm days of Spring, understanding *what* should have happened but didn't, etc. Strevens (2013) adds understanding a phenomenon *with* a theory, model, or by applying a method, to the list. It would seem there are *many* varieties of understanding.

However, as Baumberger et al. also remark, this distinction does not run very deep. For instance, understanding how a certain clock works (which would count as interrogative understanding) is the same as understanding the mode of operation of the clock (which would count as objectual understanding) (p. 30) And the same applies to understanding with a theory: if what is understood is how a pendulum clock works, then the mechanics of linear oscillations applies, together with the mathematics it relies on. So I agree with Baumberger et al.'s ensuing comment: “language does not dictate a useful distinction of different types of understanding.” The unity of intellectual understanding isn't dismissed by mere inspection of surface grammar.

3. Attribution

One might think that what we mean by “understanding” is dictated more by norms of attribution than by whatever complex cognitive state or episode is being attributed. If so, perhaps nothing independent from our attribution may exist in order to make that attribution *true*, as opposed to merely felicitous in context.⁴ In this section, I address this objection by clarifying how the attribution of understanding relates to the metaphysics of understanding thus attributed.

I will appeal to a few distinctions concerning intellectual understanding in order to defuse the objection. First, Humphreys (2000) distinguishes between primary understanding, which concerns how one discovers something for oneself, and secondary understanding. Secondary understanding is imparted understanding, where the pupil or the interlocutor learns, by testimony and perhaps further guided practice, a standardized understanding. Or perhaps the pupil merely steals the trade from one's master. Primary and secondary understanding require different abilities. Being a good communicator, for instance, matters little for primary understanding, but secondary understanding cannot obtain in its absence. In this dissertation, I consider *primary* understanding alone. Without it, after all, there would be nothing to impart.⁵

Second, Ziff (1972) distinguishes genuine understanding from misunderstanding and lack of understanding. One who thinks axioms of set theory are proper articles of faith misunderstands their roles, whereas the grade school student lacks any understanding of the axioms. Both misunderstanding

⁴ Goldfarb (1992) and Cooper (1994) advance this objection.

⁵ It is an interesting question – going beyond my project here – how my conclusions could be extended to secondary understanding.

and lack of understanding are epistemic drawbacks, in different ways.⁶ In this dissertation, I am concerned exclusively with *genuine* primary intellectual understanding.⁷

Third, Kelp (2015) distinguishes between minimal understanding, maximal understanding, and refining one's understanding. Minimal understanding is the least degree of understanding. Improved understanding goes hand in hand with improved conception of the phenomenon inquired into. And understanding is full when it excels.

The point of exploring this distinction is that, plausibly, both attributions of minimal and maximal understanding are primarily guided by pragmatic factors rather than governed by epistemic criteria. Whereas refining understanding, running the gamut from minimal to maximal understanding, seems to be the proper locus of inquiry into the metaphysics of understanding, having some freedom from pragmatic pressures concerning how we attribute understanding. Here is why.

What is it to understand something *minimally*? It seems prudent not to legislate on the matter. Many tests and standards aim to characterize what passes for minimal understanding in different research traditions. For instance, when do we pass from giving a rich enough description of Franz Ferdinand's assassination to the sketch of a satisfactory explanation of why that event occurred? Are we part of a roundtable of 20th century historians, or attending a high school history test? Context is largely going to

⁶ What genuinely understanding the axioms of set theory consists in is controversial, but we can generally tell when one has achieved such an understanding, regardless.

⁷ I remain neutral about the prospect of extending the conclusions I reach to instances of misunderstanding, in its various forms. In particular, I am neutral on whether conscious experiences that accompany understanding may have the same phenomenal character regardless of whether they accompany genuine understanding or a conception by which we partly misunderstand what we conceive. I return, in Chapter III, to how we should conceive the relationship between conscious experiences and how understanding may occur to us.

fix what passes for minimal understanding of the archduke's death. There seems to be no definitive answer to what counts as minimal *understanding* – as opposed, perhaps, to minimal belief or knowledge.

Full understanding, in the proper sense, might be unattainable, for various skeptical worries always linger. To say, with Diogenes Laertios, that only the gods fully grasp the nature of reality is a way of saying that us mortals fall short of that standard. That doesn't prevent full understanding from being an epistemic ideal which improved partial understanding aims to fulfill.

Instead, when we *say* one has a full understanding of something, we likely mean something *honorific*: they are experts, or achieved something of great value or whose achievement was difficult. (After rehearsing Gödel's proof of the incompleteness of arithmetic several times, one finally fully understands how to carry it out, so to speak.)⁸

So both attributions of full understanding and of minimal understanding seem to be responsive to context-sensitive factors rather than to the satisfaction of genuine epistemic norms. However, the *improvement* of that understanding, and our ability to subsequently exploit that understanding once already acquired, are part and parcel of our higher cognitive activity. Our thinking then seems to play an ineliminable role in *explaining* why attributions of understanding are warranted. In this dissertation, I will be primarily concerned with the central phenomenon of improving the conceptions by which we understand, or using them to properly conceive of the phenomenon understood. As Kelp writes:

I take degrees of understanding to be more fundamental than outright understanding in the

⁸ Kvanvig (2003) discusses instances of honorific attributions of understanding without presupposing that most attributions of understanding are honorific. Instead, Kvanvig offers a robust epistemic description of understanding in terms of the coherence and systematicity of our conceptions. I follow Kvanvig in focusing on instances of understanding wherein our life of thought is that primarily in virtue of which we may be attributed understanding.

sense that we first need an account of degrees of understanding in order to be able to give an account of outright understanding. (Kelp 2015, p. 3809)

However, if refining our understanding is the basic cognitive phenomenon of interest, how does it relate to outright, unqualified attributions of understanding? Kelp answers:

My suggestion is to incorporate the contextualist idea in terms of a contextually variable threshold of distance from maximal understanding. If one is close enough to fully comprehensive and maximally well-connected knowledge of a given phenomenon to surpass the threshold in a given context, then one can truly be attributed understanding of the phenomenon in that context. (Kelp 2015, p. 3813)

Kelp is here mixing a contextualist idea with his own view of maximal understanding, viz. “fully comprehensive and maximally well-connected knowledge.” I return in Section 5 below to the metaphysics of understanding, and now only approach the issue of how outright *attribution* of understanding is, in principle, achieved on the basis of degrees of understanding. And Kelp’s thought, that I endorse, seems to be the following. Relative to a given context, there will also be a threshold in the improvement of understanding that counts as good *enough*: as satisfactory understanding. Once that threshold is met,⁹ attributions of understanding become warranted. Good enough understanding warrants outright attribution relative to a contextual threshold.

These distinctions help reply to the objection this section started with. I don’t explore secondary understanding, misunderstanding or lack of understanding. And I grant that what counts as minimal or maximal understanding is to be settled by our practices of attribution. For refining understanding, when

⁹ Or there is a consensus that a threshold is met, whether that assessment is accurate or not.

outright attribution is warranted is set by a contextual threshold. With all these concessions, improved understanding is still a genuine cognitive phenomenon that needs elucidation.¹⁰ It is a question pertaining to the metaphysics of improving understanding, so conceived, that this dissertation will be concerned with. I will articulate and explore answers to the question of how phenomenal and epistemic aspects of (improved) understanding are related to each other.¹¹

4. Variety

Does the variety of understandings prevent a general account of understanding? I examine the question in this section.

¹⁰ It is important to stay neutral on what exactly constitutes a specification of the context in which we understand a phenomenon. For, if contexts were specified as necessarily being ones of inquiry (for example), then it would seem that the extent to which one understands would turn out to be dependent on the goals of inquiry. This is a problem because at issue are degrees of understanding – the state of mind attributed –, not degrees of warrant in attributing that understanding to a subject. In Section 5, I list a few epistemic norms applicable to states of mind if those states should count as states of understanding. Degrees of understanding can then be thought of in terms of how well we satisfy those epistemic norms. There are no uncontroversial ways to measure how well we are able to explain a phenomenon, or how well we exercise a given set of cognitive skills in theorizing a phenomenon. Nor is there any context-transcendent way to trade off various epistemic norms against each other. Irrespective, we all recognize, intuitively, that understanding something admits of more or less, just as our explanations are better or worse, and our thinking more or less skilled. This is so even if we cannot provide a weighted formula for how partial understanding balances the satisfaction of distinct epistemic norms in context.

¹¹ The qualifier “improved” will be presupposed in my discussion of understanding in what follows; I explicitly return to the issue in Chapter III, where I distinguish coming to understand something for the first time from exploiting an understanding one already possesses. However, it is clear that these modes of understanding may *co-occur*, for we can refine our understanding of something as we exploit it on a specific occasion. Nothing I claim hangs on appreciating this fact, though the view I advance in Chapter III does justice to it, and is more plausible for doing so.

In asking what intellectual understanding consists in, we may consider identifying conditions that are necessary for understanding, sufficient for understanding, or both. What suffices for understanding weather patterns differs so much from what suffices for understanding a religious way of life that no sufficient conditions for understanding can fruitfully be theorized. The project of giving a philosophical account of understanding has to identify necessary conditions for understanding.¹²

Variety in what we understand undermines accounts that purport to offer jointly sufficient conditions for what it is to understand in general. However, variety is no ultimate threat to identifying *necessary* conditions for understanding. Domain-specific variety is consistent with the possibility of interesting metaphysical generalizations concerning intellectual understanding. Identifying necessary conditions doesn't need to amount to a fully fleshed out philosophical account of understanding, even though it is an ingredient of any such account. (In this dissertation, I don't aim to identify *any* necessary conditions for understanding. Rather, my aim is to consider presumptive necessary conditions such as undergoing conscious experiences of understanding and satisfying a variety of epistemic norms, and ask how those presumptive necessary conditions may be related to one another.)

We may now offer a partial rejoinder to Goldfarb's challenge, in an attempt to clarify it. If Goldfarb claims that we cannot move from attributing understanding to the metaphysics of the states of

¹² To switch the example, what suffices to achieve a game-theoretical understanding of the nuclear arms' race differs from what suffices in order to possess a legal historian's understanding of *Griswold v. Connecticut*. This is so not merely because the subject-matters are distinct. For we could have an understanding of the international law surrounding the Cuban missile crisis. And we might, perhaps, also come to have a game-theoretical understanding of the progress of civil rights and liberties in the second half of the 20th century in the United States. Rather, what suffices to make our *conception* of any of these things count as understanding them likely differs from one case to another. And conceptions may be identified in terms of their subject-matters, but they may also be identified in terms of the attitudes we take in having them, or the strength with which we hold those attitudes.

mind thus attributed, then Section 3 sketches an answer. If he claims that variety in understanding (however attributed) prevents us from inquiring into the metaphysics of understanding in full generality, or somehow makes such an inquiry futile, this section sketches an answer. Doubtless, variety in the objects of our understanding is important; I return to this topic in the general conclusion to the dissertation. However, to categorically claim that there is *nothing* interesting in common to how we understand these objects seems on a par with the categorical claim that there *are* such commonalities, in that any evidence we may adduce in its favor only provides partial support. The truth, as always, lies in the details of each particular case, leaving it open whether we may fruitfully generalize across cases.

5. Decision and reality in ascertaining understanding

In this section, I point to a problem besetting theoretical views of understanding, and show how I approach it in the dissertation.

Identifying necessary conditions to be satisfied by states of mind so those states could count as states of understanding is a metaphysical project in Moravcsik's sense:

The constitutive question; i.e., what is understanding, is distinct from the epistemological issue of how we know when we ourselves or someone else has achieved understanding. (Moravcsik 1979, p. 209)

My dissertation pursues the metaphysics of understanding within that tradition.

However, when we consider satisfying *which* epistemic norms makes states of mind count as states of understanding, we meet with a problem. For almost any interesting epistemic norm – truth,

coherence, explanatoriness, being acquired by the manifestation of cognitive skills, etc. – we seem to find it intuitive that there are counterexamples.¹³

This makes it seem as though we need to *choose* which necessary conditions to impose, and to *stipulate* that only states meeting those conditions count as states of understanding, doing justice to Cooper's call to philosophically legislate:

It is the task of the normative epistemology of understanding to adjudicate between competing conceptions of understanding. (Cooper 1994, p. 15)

Here is a dilemma, starting from Moravcsik and Cooper's opposing points of view. Are we *discovering* or *inventing* what understanding is? On the one hand, if we are only discovering what understanding is, then how are we to deal with the fact that, for each epistemic norm we may consider, counterexamples seem to exist? For they would undermine metaphysical views (such as Moravcsik's) according to which meeting those norms is necessary for understanding. On the other hand, suppose we followed Cooper in thinking that we may *stipulate* which epistemic norms count for understanding. Contra Cooper, no armchair stipulation can conjure understanding into existence. Either horn of this dilemma is unappealing.

My way out of the dilemma is a *disjunctive* strategy. What makes a state of mind count as one of understanding need be no unique and logically simple criterion. Instead, understanding is *multi-dimensional*.¹⁴ I'll apply the disjunctive strategy twice: once in the epistemology of understanding, and

¹³ One may make the logical point that being *F* is a necessary and sufficient condition for being *F*. *Ditto* for being either *F* or *G*, but not *G*. (And so on for equivalent conditions.) I agree; but these conditions don't *illuminate* what understanding is. For instance, if a detective is trying to understand how a crime was committed by distinguishing necessary conditions from happenstance in its occurrence, the logical point just made will not advance her criminal investigation.

¹⁴ The existence of many dimensions of epistemic evaluation makes possible trade-offs rational. Failure to satisfy some

once in its phenomenology. In the epistemology of understanding I'll sketch in Chapter IV, I start from de Regt's (2004) view that we can only understand a phenomenon if we are able to explain it by means of a theory or model that we know how to exploit so as to extract the salient explanation of that phenomenon in the context of our inquiry. De Regt might be read as imposing a rich (and hard to satisfy) necessary condition for what should count as understanding.¹⁵ But now let's apply our disjunctive strategy. On the epistemology of understanding I offer, we may only understand a phenomenon if *either* we are able to explain it, *or* if our understanding was acquired and sustained by our cognitive skills, *or* if our understanding conceives of that phenomenon reliant upon theories and models we have of it. This states a necessary condition that has as disjuncts criteria which de Regt's view had as conjuncts.

In the phenomenology of understanding, I apply the same strategy in Chapter III. A variety of authors have identified intuitions (or intellectual seemings) as experiences of understanding. Many others have considered insights (or “Aha!” moments) as experiences of understanding. And others have noted the fluency and familiarity with which we intervene to control a phenomenon – or to solve a problem involving it – when we genuinely understand that phenomenon.¹⁶ If understanding is consciously manifested, it would seem that these are three competing descriptions of what experiences of understanding are all like. However, if we apply the disjunctive strategy, we can state that, necessarily, if understanding is consciously manifested, it is manifested *either* in an intuition, *or* in an insight, *or* in an

criteria may be offset by optimally satisfying other partial criteria in context. Closeness to the truth may give way to simplicity. Coherence may be slighted to favor fidelity to data. And so on, for other sets of possibly conflicting epistemic norms.

¹⁵ De Regt addresses scientific understanding alone, but I generalize his view. The generalization is harmless when coupled with the disjunctive strategy I mention in the text.

¹⁶ E.g., Bengson (2015), Lynch (2017), Railton (2009).

experience of fluency.

The disjunctive strategy just sketched, and applied in Chapters III and IV below, allows us to avoid both Cooper and Moravcsik's positions. *Contra* Cooper, we don't need to cherry-pick our preferred epistemic norms of understanding and stipulate that only they count. And, *contra* Moravcsik, necessary conditions for a state of mind to count as one of understanding aren't something as easily stated as, say, the intellectual perception of real essences (a quasi-Aristotelian view Moravcsik sketches). These conditions are, rather, highly disjunctive and not easily specified – as we may have expected for a state of mind such as understanding.

6. The epistemology of understanding

In the last three sections of this Introduction, I summarize the dissertation to follow. The dissertation approaches the question of how epistemic and phenomenal aspects of understanding are related. In order to address that question, preliminary accounts of understanding's epistemic and phenomenal aspects are needed. In Chapters I-II, I discuss epistemic aspects of understanding. In Chapter III, I discuss phenomenal aspects of understanding. Chapter IV then raises the question of how these aspects are related.¹⁷

¹⁷ The starting point of my inquiry is van Fraassen (1999), who suggests that religious revelations and conversions to a new scientific paradigm share important phenomenal and epistemic properties. However, my project improves over van Fraassen's in a number of ways: being concerned with understanding rather than other states of mind; being concerned with understanding in full generality rather than only in scientific paradigm changes; and in asking what kinds of experiences may constitute understanding rather than merely accompany it.

In Chapter I, I explore how to formulate a virtue epistemology of understanding, faced with how widespread and far reaching our cognitive biases are. On Zagzebski's (2001) view, the conceptions underwriting how we understand something – for instance, the position on a chessboard – gain their positive epistemic status by having been formed in the right way, where proper formation involves the manifestation of intellectual virtues or cognitive skills (like open-mindedness or imaginativeness). Chapter I considers whether cognitive biases undermine the positive epistemic status of understanding, and supports an alternative to reflection-based conceptions of expertise. Carter and Pritchard (2016) raise a skeptical objection to virtue-epistemological approaches to understanding relying on the ubiquity of biases. I first argue that their objection is best fleshed out in terms of cognitive biases like unguarded uses of the availability heuristic, the hindsight bias, or confirmation biases. I then defend the grounding of understanding in intellectual virtues and skills from Carter and Pritchard's objection. Even if Carter and Pritchard's objection succeeds against Zagzebski's specific view, the underlying reason is that both view and objection assume that understanding something is necessarily accessible to us in reflection. In contrast, I will argue that we are cognitively responsible for understanding whenever that understanding is primarily formed by manifesting our expertise, and that we can understand without being required to be able to reflect upon our own understanding. So, even if biases are opaque to reflective scrutiny, their hiddenness does not, on its own, undermine our cognitive ownership over the conceptions by which we understand.

In Chapter II, I address how views of understanding in general fare when we consider instances of mathematical understanding. According to Grimm (2014), we only understand a phenomenon if we are able to explain its occurrence, and we can only explain a phenomenon if we know what other phenomena it depends on. We identify dependencies, Grimm claims, according to how we answer “What

if things had been different?” questions. I argue that Grimm’s view of understanding meets with mathematical counterexamples. I offer two such examples: a simple arithmetical argument and a set-theoretic axiom. These are counterexamples to Grimm’s view because, in mathematics, things couldn’t have been different.

I consider three replies Grimm may make, according to how we might interpret “What if things had been different?” questions, and I argue they do not succeed in rescuing Grimm’s view from counterexamples. On the first reply, answering “What if things had been different?” questions enables us to grasp the modal profile of what is understood. On the second reply, answering such questions enables us to know what grounds what. And, on the third reply, “What if things had been different?” questions consider not necessarily differences in phenomena so much as differences in how we model those phenomena. While each approach may seem promising in its own right, I will argue that they cannot aid Grimm’s view escape counterexamples.

The falsity of Grimm’s view doesn’t imply that understanding a phenomenon couldn’t often involve being able to explain its occurrence, nor that such an explanation doesn’t often seize upon what that phenomenon depends on. Rather, I argue for a disjunction. Either, firstly, “What if things had been different?” questions can’t be asked in the same way regardless of what phenomena are being understood. Or, secondly, understanding doesn’t always involve the ability to explain what is understood. Or, thirdly, the explanation in question doesn’t always seize upon what the phenomenon understood depends on. I stay neutral about which of these possibilities obtain. This disjunction affords preserving the traditional view that understanding and explanation are closely related.

Summing up, Chapters I-II reject an objection to a widespread view of understanding (responsibilist virtue theory), while raising objections to another widespread view (Grimm’s view).

Jointly, these two points support a form of pluralism about the epistemic norms of understanding: how we identify what makes understanding possess the positive epistemic status it does may be approached in a variety of ways, and the success or failure of such views isn't guaranteed, but depends on their details, and on the arguments proposed for or against them. I use the label "intellectual understanding" precisely in order to stay neutral concerning what the ultimate epistemology of understanding in general might be, if any such view can ever be adequately supported.

7. The phenomenology of understanding

My dissertation pursues the question of how epistemic and phenomenal aspects of understanding are related. Chapter III discusses the phenomenology of understanding. I start from two recent views, those of Bourget (2017) and Lynch (2017). For Bourget, understanding *P* partly consists in grasping *P*, and the occurrence of its grasp, in turn, consists in having a conscious experience with *P* as content (though, if memory fails us, for instance, that content is only implicit in our experience and may be articulated in dialogue). And, for Lynch, we identify the state of mind that qualifies as understanding (because it enables us to explain what is understood, or to solve problems that appeal to what we understand) in terms of the initial act of coming to understand. Coming to understand, Lynch argues, is a creative act, and its creative aspect partly consists in having conscious insights, whose felt novelty and surprise indicate to us that we may have genuinely understood something.

Although I argue that both Bourget and Lynch's views are implausible, both views are starting points for proposing the phenomenology of understanding I favor. On the view I advocate, Bourget and Lynch's claims aren't rejected. Rather, I stay neutral on whether such claims are true. Avoiding

commitments that seem implausible is a *pro tanto* reason to endorse the phenomenology of understanding I propose.

On the phenomenology of understanding I propose (in view of addressing the question of how phenomenal and epistemic aspects of understanding are related to each other), any instance of understanding occurs in one of three modes. Each mode may be consciously experienced.¹⁸ Moments of coming to understand are conscious in insights, or “Aha!” experiences. Exploiting an understanding one already possesses in a discursively articulated way is conscious in intuitions. And exploiting an understanding one already possesses but which one cannot discursively articulate (an otherwise tacit understanding) is manifest in experiences of fluency.

Consider the claim that some conscious experiences – like insights, intuitions, and experiences of fluency – more closely align with modes of understanding than *other* experiences – say, relief at having finally understood something, or enthusiasm at bringing to bear a theory one understands well to a previously mysterious phenomenon. This claim raises the question of precisely how to draw the distinction between experiences that the occurrence of understanding might possibly consist in or is typically caused by, and experiences that accompany the occurrence of understanding by happenstance alone. So as not to prejudge which conscious experiences might more closely align with understanding, that understanding has to be described in terms of its positive epistemic status. And this brings us back to the question of how phenomenal and epistemic aspects of understanding are related, a question Chapter IV addresses.

¹⁸ I don't presuppose that understanding something need always be accompanied by conscious experience. But the possibility of unconscious understanding should not prevent us from inquiring into how the phenomenal character of experiences associated with understanding may contribute to the epistemic properties of that understanding.

8. How the epistemology and the phenomenology of understanding are related

When a chess master looks at a chessboard, she grasps the patterns on the board, playing fluently, having the right intuitions about which plans are feasible, and sometimes experiencing insights concerning combinations that may destabilize the opponent's position. The player's understanding of a chess position typically has both phenomenal and epistemic aspects. But do those aspects merely co-occur, or is there any discernible relationship between the epistemic grasp that the player has of the position on the board and what she consciously experiences?

Lest all conscious experiences should meet the epistemic norms of understanding merely by coincidence if and when they do so, it would seem that we also need to answer what I will call "Wittgenstein's question:" the question of *why* phenomenal and epistemic aspects of understanding are related the way they are supposed to be (in each specific view of understanding). In Chapter IV, I'll argue that we lack the evidence needed to adequately support the claim that what it is like to experience understanding is related to meeting the epistemic norms of understanding by a single, metaphysically necessary, relationship. We may ultimately be unable to satisfactorily answer Wittgenstein's question.

After I sketch an attempt to answer Wittgenstein's question made by Bengson (2015) and raise an objection to his view, I'll support the conclusion that Wittgenstein's question may ultimately not be given a satisfactory answer by considering a view that isn't held by any one author, but seems to me to be the most promising in addressing Wittgenstein's question. I reason hypothetically: if the view that seems most promising in tackling that question doesn't succeed in giving it a satisfactory answer, then no such answer is likely to be forthcoming.

On the view I find most promising for addressing Wittgenstein's question, insights, intuitions, and experiences of fluency are conscious experiences that have been systematically linked to the occurrence of understanding (in agreement with the view proposed in Chapter III). And the use of models, the ability to explain target phenomena, and manifesting cognitive skills are *possible* epistemic norms governing understanding (in agreement with the pluralism concerning what norms apply to understanding that is supported by Chapters I and II).

How do these particular conscious experiences contribute to meeting these particular epistemic norms of understanding? The hypothesis I offer is that we understand by attending. That is, different experiences of understanding meet distinct epistemic norms partly because, in each type of experience, we consciously attend to the phenomena understood in a manner specific to that kind of experience.

Here, in a nutshell, is how phenomenal and epistemic aspects of understanding might answer each other – only as a proof of principle. In insights, we shift our patterns of attending to what is understood so that, while we couldn't explain it before having the insight, we are enabled to explain what we understand by undergoing that “Aha!” experience. In intuitions, we focus our attention to the phenomenon intuited, and the attention we focus is framed in terms of the models we have of those target phenomena we attend to. And experiences of fluency that show our understanding of something also manifest the cognitive skills at work in acquiring and sustaining that understanding.

These correlations, however, leave open the question of whether understanding gains its positive epistemic status *in virtue of* being experienced (when such experiences occur). In addressing this latter question, I argue that the evidence available concerning how attention is related to conscious experiences and, respectively, to meeting epistemic norms, can at best support claims made with the modal strength of natural necessity. And that leaves open the question of whether that natural necessity is grounded in a

metaphysical relation of causation or constitution. Conceived so as to require answers advanced with the modal strength of metaphysical necessity, Wittgenstein's question can't be given a fully satisfactory answer. That question pertains to *why* phenomenal and epistemic aspects of understanding may relate as they supposedly do. Failing to answer that question supports the view that we have no decisive *rationale* for thinking that epistemic and phenomenal aspects of understanding do relate in any metaphysically salient general way.

Chapter I

Intellectual virtues and biased understanding

Biases affect much of our epistemic lives; do they affect how we understand things? For Zagzebski (2001), we only understand something when we manifest intellectual virtues or skills. Relying on how widespread biases are, Carter and Pritchard (2016) raise a skeptical objection to understanding so conceived. It runs as follows: Most of us seem to understand many things. We genuinely understand only when we manifest intellectual virtues or skills, and are cognitively responsible for so doing. Yet much of what we seem to understand consists in conceptions whose formation could have easily been due to biases instead. And the work of biases is opaque to reflection. If conceptions constituting how we understand things could have easily been due to biases, then we're not cognitively responsible for them because we can't reflectively endorse them. So we are mistaken in thinking we genuinely understand most of the time. I will defend the grounding of understanding in intellectual virtues and skills from Carter and Pritchard's objection. We are cognitively responsible for understanding when we manifest our expertise. We can do so, I will argue, without being required to reflectively endorse our own understanding.

1. Understanding

We use the word “understanding” in many ways. Honorific uses aside, we may depict several varieties of understanding; let's illustrate them. When we understand what someone says, perhaps we only know what they mean. When we understand what we see, perhaps that only comes down to segmenting the visual scene by applying concepts. When we understand what a friend is going through, perhaps that only calls for empathy. When we understand how to ride a bike, perhaps that only involves the ability to do it and familiarity with having done it in the past. Linguistic, perceptual, emotional, and practical understanding doubtless exist.

However, I will focus on intellectual understanding, following Zagzebski (2001), who is the main proponent of a virtue epistemology of understanding. I do so in order to consider and reject a skeptical challenge to virtue epistemologies of understanding. The challenge, raised by Carter and Pritchard (2016), relies on the fact that biases are ubiquitous and pervasive in thinking. In the rest of this section and the next, I sketch the relevant details of Zagzebski's virtue epistemology of understanding. In Section 3-5, I formulate Carter and Pritchard's skeptical challenge starting from the pervasiveness of biases and the risks they pose to thinking. Sections 6-7 articulate my response to the skeptical challenge on behalf of a virtue epistemology of understanding, which appeals to a notion of responsibility that doesn't presuppose we're able to reflectively endorse the conceptions by which we understand.

What is understanding?

Understanding is a cognitive state that arises from *technê* ... The person who has mastered a *technê* understands the nature of the product of the *technê* and is able to explain it. (Zagzebski 2001, p. 240)

For Zagzebski, to understand something manifests skill. One's understanding is "the product of the *technê*" – and one has "to be able to explain it." For instance, what is it to understand the position on a chessboard? Most players can, on inspection, identify the state of the game they are playing, forming – to the best of their abilities – conceptions about, e.g., how developed their pieces are, what weaknesses the opponent's position has, and how those can be exploited. Conceptions of this sort, which constitute how players understand a position, are available for them to use in explaining that position.¹⁹ This view of understanding is in the background of Zagzebski's virtue epistemology, to which I now turn.

2. Virtues and skills

Zagzebski writes: "Understanding is an epistemic state that arises from *technê*" (p. 240). *Technai* are virtues or skills. Which ones are at play in understanding varies from case to case. Understanding some phenomena may require open-mindedness to see alternatives where others don't. And understanding other phenomena may require a vivid imagination.

Zagzebski's virtue epistemology is best construed as addressing four aspects of understanding:

¹⁹ Zagzebski has a richer conception of understanding, writing: "Understanding involves seeing how the parts of that body of knowledge fit together, where the fitting together is not itself propositional in form" (Zagzebski 2001, p. 243). The resulting conception isn't clear, for three reasons. First, Zagzebski agrees that, although skill-based understanding isn't *propositional* on the face of it, that doesn't imply one cannot propositionally specify the conceptions by which one understands; so one shouldn't be sanguine about non-propositionality. Second, it's unclear whether the requirement of *coherence* (or fitting-together) applies to "parts of that body of knowledge" or to our conscious experience of it. Third, the quote suggests that only parts of a "body of knowledge" may cohere to form an understanding. Yet she also writes that "understanding does not always build on a base of knowledge" (2001, pp. 243-244).

justification, formation, excellence and responsibility. The conception by which we understand something is justified because it was formed in the right way: its formation, in the initial episode of coming to understand, must have primarily manifested intellectual virtues or skills appropriate to the context of inquiry. The manifestation of virtues and skills explains what is excellent in cognition, aiming at optimally meeting epistemic norms such as accuracy, coherence or explanatoriness.

Other epistemologies of understanding can also account for the formation, justification and excellence of the conception by which we understand something. Consider reliabilism. As Zagzebski (2001, p. 238) points out, virtues and skills are reliable: those who understand possess varying degrees of expertise, and an expert is “the person who is the most reliable source of knowledge” in her field of inquiry. The reliability of virtues and skills explains why the conceptions they help form are mostly true. Or consider coherentism. “Understanding involves seeing how the parts of that body of knowledge fit together” (Zagzebski 2001, p. 243). Presumably, it is the conceptions of experts that are most coherent in their domains of expertise.

What's novel in Zagzebski's virtue epistemology is a focus on epistemic agents. In understanding, we don't merely house special states of mind. Rather, we are responsible for how we understand things; we author our understanding inasmuch as *our* virtues and skills explain how we came to conceive what we understand.

With this sketch of a virtue epistemology of understanding in the background, we can now consider Carter and Pritchard's skeptical challenge to it.

3. Biases

Carter and Pritchard's skeptical challenge relies on the ubiquity and importance of biases in everyday thought. To set the stage for their challenge, this section asks what biases are. Carter and Pritchard say:

A bias, in the most general sense, is a disposition, implicit or explicit, to reach a particular kind of conclusion or outcome – in the kind of case we're interested in, the outcomes will be representational. Call these *cognitive* biases. (Carter and Pritchard 2016, p. 273)

This description doesn't lead far. Any disposition is “implicit or explicit.” And any disposition is manifested in “a particular kind of conclusion or outcome.” So far, all we have to go on is that cognitive biases are dispositions to produce (presumably cognitive) representations. This doesn't say why biases threaten understanding. To properly appreciate the skeptical challenge that biases occasion, it pays to improve on Carter and Pritchard's characterization of biases. Two options seem open for what might be wrong with biases. Either the representation produced is defective: biases help form conceptions which are systematically *erroneous*. Or the way in which representations are produced is defective: biases taint the *etiology* of the conceptions they help form. Let's take them in turn to see if fleshing them out can occasion a skeptical challenge.²⁰

²⁰ Carter and Pritchard (2016, p. 273) follow Saul (2013) in characterizing any bias whose outcome is representational as cognitive. This effaces the distinction between biases that do and those that don't often trigger *affect*. The latter include biases in categorizing black men as more dangerous, women as less competent, or the poor as lazier. Carter, Pritchard, and Saul seem to assume that both sorts of biases pose a similar epistemological threat – to knowledge, understanding, or our capacity as inquirers – and they sensibly focus on biases with a higher affective and practical impact. However,

Kahneman (2002, pp. 3-4) voices the often met view that biases typically produce inaccurate representations: “Systematic errors are known as biases, and they recur predictably in particular circumstances.”²¹ However, the defective-representation construal faces a problem: it’s unclear why biases should systematically depart from the truth in all circumstances. We often have to reach decisions fast, and in conditions of uncertainty. If fast and frugal heuristics are effective in advancing our adaptive goals, it might be misguided to expect more from them (Chase, Hertwig and Gigerenzer 1998).²²

A skeptical challenge might more readily be mounted based on how biases feature in the *etiology* of our conceptions. In virtue epistemology, conceptions constituting our understanding are justified by having been formed in the right way – primarily, through manifesting our intellectual virtues or skills. Why would biases playing a role in how our understanding is formed be epistemically vicious?

understanding rarely evokes a powerful affective response in us. Because I believe they may more readily sustain a skeptical challenge, I focus on cognitive biases that seem of wider applicability, e.g. overconfidence. (Morewedge and Kahneman 2010 call these “judgment biases.”) And I illustrate these biases with chess scenarios, where social and affective consequences are comparatively negligible.

²¹ On the defective-representation view, errors are no accident; we can give a just-so explanation for biases. Conceptions constituting our understanding are arrived at *via* reasoning, and that takes time and effort. Often, under the pressure of time, we apply shortcuts in reasoning – heuristics – that trade speed for accuracy. We have evolved some useful heuristics: often we need to make judgment calls in the blink of an eye, and don’t have time to reason things through. Applying heuristics may fall short of perfect accuracy, as characterized by the laws of classical probability theory. Forming conceptions fast rather than accurately may then introduce error.

²² Perhaps the appearance that biases produce inaccurate representations is only due to unrealistic benchmarks for what should count as accurate human reasoning. To remedy this, one may, for instance, follow Fennell and Baddeley (2012) in mixing prior probabilities of ignorance (e.g., 0.5) and prior probabilities learned through experience. The mix would be weighted by how often situations encountered were familiar to the subject (needing priors of experience) and how often situations encountered were unfamiliar (needing priors of ignorance). I remain neutral concerning the proper formal representation of biases. Indeed, perhaps pluralism is in order, so that what count as biases by some standards may be acceptable cognitive procedures by other standards. My goal is more limited: to explore if biases may support a particular skeptical challenge to virtue-based understanding.

Morewedge and Kahneman (2010, p. 435) say:

It is often useful to think of judgments as a weighted combination of items of information... In this scheme, judgment biases can always be described as an overweighting of some aspects of the information and underweighting or neglect of others, relative to a criterion of accuracy or logical consistency... In this fashion, the principles of associative activation help explain biases of judgment.

They see the most common cognitive biases as side-effects of associative semantic memory. Abstracting away from the details of their hypothesis, the threat biases pose seems to be the following. How activation spreads within a semantic network is not up to the agent. It is an informational and biological *datum*. So, even when conceptions which partly result from biases are accurate, it's far from clear that we are responsible for what we believe: because we didn't have a say on how those conceptions were formed, and can't effectively control them. This worry about how we control our own thoughts – how they are, in a sense, *ours* – seems principled. The worry applies most poignantly to those thoughts we value most, the ones by which we understand the world around us. I'll now argue it is this kind of worry that motivates Carter and Pritchard's skepticism.

4. Carter and Pritchard's challenge

In this section, I sketch Carter and Pritchard's bias-driven skepticism.²³ They write:

²³ Saul (2013) supports bias-driven skeptical conclusions about knowledge. Carter and Pritchard (2016, p. 288), however, think that virtue epistemology doesn't apply to knowledge. (To give just one example of why not: We often gain knowledge

On the supposition that the individual's cognitive success is not primarily explained by her exercise of cognitive ability, but also partly down to an unconscious bias, is one able to enjoy the kind of cognitive ownership of this fact that is characteristic of understanding? We suggest not. (p. 287)

Grant that understanding manifests virtues and skills. However, most of us are often subject to biases,²⁴

by testimony. Even when listeners are discerning, knowing who to trust, knowledge by testimony isn't primarily due to their virtues.) However, Carter and Pritchard think virtue epistemology applies to understanding. So they correspondingly consider skeptical theses based on how biases would undermine understanding. Carter and Pritchard also consider Alfano's (2012) situationist challenge to virtue epistemology as raising a similar epistemic challenge to that of Saul (2013). However, to the extent that biases are dispositional, they can be characterized across a range of different situations, so situation specificity raises different concerns than biases. Alfano's challenge goes beyond the scope of this paper; Railton (2012) outlines a reply to situationism.

- ²⁴ It would be unrealistic to focus on those rare cases where experts *lack* bias entirely. However, it is reasonable to assume that many of the conceptions by which we understand aren't *primarily* due to bias. When could forming an understanding of something have *easily* been due to biases instead of intellectual virtues or cognitive skills? The possibility of bias is *live* if the probability of a bias-inclusive etiology of our conceptions *isn't* low. Yet the probability of a bias-inclusive etiology of the conceptions by which we understand couldn't be *too* high either. Otherwise, their virtuous etiology would be unexplainable, sheer luck. How high the probability of a biased etiology of the conceptions by which we understand, then, has to be a context-sensitive matter: neither too low (so as to make it negligible) nor too high (so as to undermine the actual operation of our virtues and skills).

Oddly, Pritchard doesn't offer a probability-based answer to the question of what "easily" means here, which is less than ideal because it doesn't make the role of context explicit. Instead, Pritchard offers an answer based on nearby possible worlds. He interprets the phrase "it could have easily been the case that *p*" to mean "*p* holds in most nearby possible worlds."

Pritchard (2014) appeals to this analysis of "easily might have been" because he gives a modal account of both epistemic luck and epistemic control. On his view, loss of our cognitive control over how we form our conceptions should be construed largely the same way as the occurrence of epistemic luck. In reply, I doubt that a modal account of epistemic agency is adequate. For why would the mere modal *correlation* between, say, my intention to think what's right and my actually thinking what's right speak to how my intention makes my cognitive behavior *intelligible*? Fridland (2014) points to problems with theorizing control, epistemic agency, and the attendant luck. I conjecture that some of these problems arise because different varieties of control (and luck: environmental, reflective, etc.) should be represented differently,

enough so that we can't sort out biases or virtues-plus-skills as primary causal factors in how we have come to understand something. Then, even if our conceptions were in fact primarily due to virtues and skills, they could have easily been primarily due to biases instead. Why is this possibility problematic for understanding? Because biases, being sub-personal cognitive routines that operate relatively automatically and implicitly, undermine our cognitive responsibility – “the kind of cognitive ownership... that is characteristic of understanding.”²⁵ As Carter and Pritchard put it:

Like cognitive achievement, understanding requires a particular kind of cognitive ownership, that isn't essential for merely knowing. But, unlike cognitive achievement, that cognitive ownership also essentially involves a reflective grip on the matter in hand, one that would require the subject to have a rational basis to exclude live-error possibilities. (p. 288)

Cognitive responsibility for understanding implies that one should have a basis, in reflection, for excluding “live-error” possibilities. What are some live-error possibilities being excluded? If the conception constituting how we understand something could have *easily* been due to biases rather than our intellectual virtues or skills, then the possibility of an epistemically vicious etiology for what we seem to understand is *live*. The problem is that, since biases often operate implicitly, we typically *don't* have a basis in reflection for what we seem to understand whenever what we seem to understand could

including in a possible-worlds framework. Offering a positive account, and modal representation, of epistemic agency go far beyond the topic of this text. Here, I only gesture toward why one can question Pritchard's construal of alternative non-actual etiologies for our conceptions, etiologies that might easily have obtained. I am grateful to Paul Humphreys for discussion about these issues.

²⁵ What is it to have *ownership* over one's understanding? We are obviously bearers of our states of mind. The question of ownership has to turn on whether we're also cognitively responsible for them, as authors.

have easily been due to biases rather than our intellectual virtues or skills. By their reflective opacity yet cognitive efficacy, biases undermine our cognitive responsibility. Since understanding implies cognitive responsibility, we often don't *genuinely* understand what we *seem to*.

Although they only aim to *motivate* a skeptical worry, I believe Carter and Pritchard's skeptical challenge is best represented in argument-form. Carter and Pritchard don't offer this argument word for word. My own reply to their skeptical challenge depends on how faithful to their intentions this argument is:

- (1) All instances of genuine understanding are primarily due to the epistemic agent's intellectual virtues and skills.
- (2) If an instance of understanding is primarily due to an epistemic agent's intellectual virtues and skills, then that agent is cognitively responsible for it.
- (3) If an epistemic agent is cognitively responsible for understanding, then that agent must be able to endorse that understanding in reflection.
- (4) Most instances of seeming to understand could have easily been due to biases (though they weren't).
- (5) If an instance of seeming to understand could have easily been due to biases (though it wasn't), then the epistemic agent couldn't endorse it in reflection.
- (6) So, most instances of seeming to understand are not instances of genuine understanding.

The premises each do important work in supporting the conclusion. (1) is key to how virtue epistemology construes understanding: as a state of mind individuated in terms of its acquisition – an acquisition the epistemic agent can be credited for. I defend understanding, construed as based on virtues

and skills, against skepticism; so I grant (1).

Including (2) purchases Carter and Pritchard a narrower target. If one had a reliabilist approach to intellectual virtues or skills (Greco 1999), instances of understanding that are luckily free from bias could be recognized to exist by denying (2). Carter and Pritchard's challenge only properly applies to responsibilist virtue epistemologies of understanding,²⁶ e.g., Zagzebski's (2001), not to reliabilist ones. Given the dialectic, we should grant (2).²⁷

(4) is an empirical premise to the effect that, given how widespread biases are, it is always possible that conceptions formed well, and constituting our understanding of something, might have easily been produced primarily by our biases rather than our intellectual virtues and skills. I, of course, grant that biases are widespread, and that they often insinuate themselves in how we think.

(5) concerns how biases interact with reflection. No one can reflectively endorse conceptions that

²⁶ Pritchard (2005, p. 190) explicitly contrasts credit due to an epistemic agent for conceptions formed by manifesting reliable competencies with the cognitive ownership characteristic of manifesting intellectual virtues we are responsible for. He writes that “the cognitive achievement is entirely at a sub-personal level, and in this sense the agent proper is not cognitively responsible for her reliably formed true beliefs at all... agent reliabilist accounts of knowledge might ensure that agents are able to take a very minimal form of cognitive responsibility for their beliefs,” viz., credit. Contrast this with Carter and Pritchard's (p. 288) view that “understanding requires a particular kind of cognitive ownership” that includes but goes beyond credit because it “essentially involves a reflective grip on the matter in hand, one that would require the subject to have a rational basis to exclude live-error possibilities.” Call this requirement “reflective endorsement;” I'll return to it in Section 5.

²⁷ This is not to say that the distinction between purely reliable competences and virtues we are responsible for isn't controversial; Baehr (2006) suggests that these two aspects may often presuppose each other. While the mental items discussed may be the same, responsibilism and reliabilism still differ on what has explanatory priority: the fact that we are responsible for the intellectual virtues in question; or the fact that we thus get at the truth. The equation “Truth is what the sage aims to believe” may be read differently according to whether one emphasizes the left-to-right or the right-to-left order of explanation; cf. Blackburn (2001).

are reflectively opaque. As biases typically operate implicitly, their absence is equally unavailable to reflective scrutiny.²⁸ Since Carter and Pritchard's challenge trades on what typically happens when biases luckily fail to be primary causal factors in how we come to understand something, we should grant (5).

I argue against skepticism about understanding. So I deny conclusion (6) that most instances in which we seem to understand something aren't instances in which we genuinely understand. The premise I deny is (3). Before arguing in Section 7, against (3), that one can be cognitively responsible without reflectively endorsing one's understanding, I'll clarify the central terms involved: the appeal to reflection (Section 5), and our responsibility for manifesting virtues or skills in understanding (Section 6).

5. Reflective endorsement

In Section 4, I called the requirement Carter and Pritchard impose on the kind of cognitive responsibility we have for understanding “reflective endorsement.” Their exact words refer to a “cognitive ownership [that] essentially involves a reflective grip on the matter in hand, one that would require the subject to have a rational basis to exclude live-error possibilities” (p. 288). Notice two key phrases: “reflective grip” and “rational basis.”

²⁸ As Wilson et al. (2002) note, conscious de-biasing in reflection is difficult but not impossible. It is, then, open to a virtue epistemologist to say that we might form some conceptions early on, under the influence of bias, and only *later* those conceptions may come to constitute a genuine understanding – later, if and when de-biasing and reflective endorsement becomes possible. One should note that reflective access and reflective effectiveness differ greatly; one may be capacious enough as to become aware of a bias one has, while unable to change the pattern of vicious thinking however much one tried. Alternative strategies like changing the environment or focusing on areas where the bias isn't operative may work better (Antony 2016). I'm grateful to an anonymous reviewer for the point.

What do they mean? One has a *reflective grip* on a matter when one can reflect on it, and one's grip consists in one's ability to reach some conclusion about that issue in reflection. In particular, when it comes to the justification of conceptions that constitute how we understand something (e.g., the position on a chessboard), one can reach, in reflection, some conclusion – however tentative and qualified – about whether one is right to have the conceptions in question or not; whether they're justified by one's own lights. One has a *rational basis* for holding the conceptions one does when, on balance, the reasons favoring one's conceptions outweigh the reasons against them.²⁹ The requirements of having a reflective grip on what one understands and having a rational basis for it are logically independent.³⁰

On Carter and Prichard's view, we count as understanding something only if the conceptions that constitute our understanding can sustain reflective scrutiny *because* we have a rational basis for them available to us in reflection. In Section 4, I called this complex requirement "reflective endorsement." It adds the reflective-grip to the rational-basis requirement.³¹ In Sections 6-7, I will argue that experts *don't*

²⁹ One could sharpen this requirement further in several ways, adding that the epistemic agent should be able to become aware *of* one or more reasons favoring what they believe, or that such reasons should be *salient* in the context of inquiry. A good menu of options of extra requirements is Owens (2000). However, believing what, on balance, one's reasons favor, suffices for having a rational basis.

³⁰ On the one hand, what one's conceptions are rationally based on may be unavailable to the epistemic agent's reflection. A rational basis for what one believes may be ascertained by one's responsiveness to reasons, or rational sensibility. A rational sensibility which doesn't always involve reflection seems to underlie the very manifestation of virtues and skills. On the other hand, having a reflective grip on one's conceptions need not be rationally based. We may, in reflection, identify different reasons, or weigh them differently, than what genuinely supports or undermines our conceptions (Dancy 2006). And our reflection may also assess not whether our conceptions are based in reasons, but whether they promote our projects or are otherwise useful.

³¹ Carter and Pritchard are concerned with understanding's having a rational basis "to exclude live-error possibilities," where a biased etiology counts as the relevant possibility envisaged. Having "a rational basis to exclude live-error possibilities" sounds ambiguous (to me, at least). For it is one thing to have reasons in favor of the conception by which I understand the position on a chess board, reasons strong enough that they – in fact, and my awareness of them regardless – exclude

always have a reflective grip on what they understand. Hence neither having a reflective grip on what is understood nor being able to reflectively endorse what is understood can be required for understanding – even when one's cognitive behavior is responsible. It follows that premise (3) in Carter and Pritchard's argument is false.

What is the relation between responsibilist virtue epistemology and a reflective endorsement requirement on understanding? For Carter and Pritchard's challenge to hit its target,³² it has to be the case that reflective endorsement, as a requirement on cognitive responsibility, somehow flows from the fact that one comes to understand by manifesting intellectual virtues or skills. For all “neo-Aristotelian”

even live-error possibilities. And it is another thing to have a second-order reason (in Dancy's 2006 terminology), ascertained in reflection, to think that first-order reasons in favor of our conceptions outweigh or override first-order reasons against our conceptions. A virtuous etiology is a reason in favor of our conceptions; the non-negligible risk of bias interference is a reason against them. This richer construal of the reflective requirement Carter and Pritchard impose on our cognitive responsibility for understanding includes, and is perhaps stronger, than both the reflective grip and the rational-basis requirement. So, if what I call “reflective endorsement” in the text is too cognitively demanding to fit all cases of understanding (as Section 6-7 below argue), then the criticism of over-demandingness will apply to Carter and Pritchard's stated requirement as well.

³² Suppose that being able to reflectively endorse your understanding were a dispositional condition coextensive with the standing conception by which you understand. *Whenever* you would use that conception, you should be able to ascertain, in reflection, that your reasons for holding that conception outweigh or override your reasons against holding that conception – if you genuinely understood what you conceive. So construed, reflective endorsement is independent from the requirement that your conception be *formed* by manifesting intellectual virtues and skills. There is a problem with construing reflective endorsement and virtue-manifestation as independent components of cognitive responsibility. The dimension of epistemic appraisal involving reflective endorsement makes *no* reference to the epistemic agent's virtues and skills at all. If so, Carter and Pritchard's would, in effect, *beg the question* against a responsibilist virtue epistemology of understanding. For a crucial epistemic aspect of how we understand and what makes us responsible for understanding – the amenability of our conceptions to reflective scrutiny – would *not* be accounted for in virtue-theoretic terms, *contra* the spirit of a virtue epistemology of understanding. To avoid begging the question, the ability to reflectively endorse our understanding has to be construed differently, as somehow flowing from the manifestation of intellectual virtues and skills.

or responsibilist virtue epistemologies, writes Pritchard,

epistemic virtues are reliable cognitive traits which also demand a certain level of reflective responsibility on the part of the agent. (Pritchard 2005, pp. 194-195)

“Reflective responsibility” meshes well with Zagzebski’s view, on which “it is impossible to understand without understanding that one understands” (Zagzebski 2001, p. 246).³³ And the reflective dimension of understanding matters for cognitive responsibility: “One of the central features of agency is self-reflectiveness” (Zagzebski 2001b, p. 152). So Carter and Pritchard’s argument seems to *succeed* against Zagzebski because she accepts (3). Yet, as Pritchard (2005, p. 195) notes, Zagzebski’s is only “a version” of virtue epistemology. In Section 6, I will argue that a responsibilist virtue epistemology of understanding is coherent even in the absence of a reflective endorsement requirement on cognitive responsibility.

6. Acts of virtue and skill

In this section, I sketch a notion of cognitive responsibility for the irreflective manifestation of our intellectual virtues and skills, *contra* (3).

³³ Taking Zagzebski at face value about this is unpromising. Compare Cargile (1970) on knowing that one knows, remarks that could easily be adapted for understanding that one understands. The logical charge of regress (how do you understand that you understand that you understand – *p*?), albeit reasonable, doesn't get a bite, Cargile suggests, because no recursive relation has been specified between the worldly (or mental) *contents* of the embedded “that”-clauses, even though generating those clauses themselves is grammatical (abiding by recursive rules governing *words*). The more pressing problem, Cargile thinks – and I agree –, is that the letter of conditions like Zagzebski's make knowledge or understanding cognitively implausible, for such conditions are too demanding.

Zagzebski's virtue epistemology of understanding presupposes that one is cognitively responsible for how one understands something (e.g., the position on a chessboard). Both I and Carter and Pritchard share that assumption. That raises the question of what cognitive responsibility for understanding is.³⁴ Watson writes:

Because many of these appraisals concern the agent's excellences and faults – or virtues and vices – as manifested in thought and action, I shall say that such judgments are made from the *aretaic perspective*. (Watson 1996, p. 231)

The sage is responsible for their deed *because* they have the virtues and skills that deed manifests. Those virtues and skills³⁵ are part of *who one is*, and their development is part of one's learning history. (Since character is usually seen as comprised of virtues, I'll speak of expertise in a domain as a catch-all for one's skills in that domain.)³⁶ Suppose you understand the moral horrendousness of a situation, or that you understand that the eradication of poverty isn't feasible. To be responsible for how you understand that is, then, to be responsible for the virtues and skills manifested in the formation of the conception by which you understand what you in fact do. To be responsible for those virtues and skills is for them to be

³⁴ One unsuitable view is doxastic voluntarism. Surely we don't come to understand things the way we do because *we wish it*. From a virtue-theoretical perspective, to think that cognitive responsibility for understanding has to amount to doxastic voluntarism is to conflate *aretaic* and act-relative perspectives in attributing responsibility. For how these perspectives differ, see Watson (1996) and Zagzebski (1996).

³⁵ Watson (1996, p. 244) argues against the generalization of this notion of responsibility from virtues to skills. To reply: virtues do differ from skills in important respects. But that falls short of showing that the notion of responsibility applicable to acts of virtue doesn't apply to acts of skill. If Hans writes a complex logical proof with ease and skill at calculation, we may properly say he's responsible for it: not for his scribbles on paper but for the proof – precisely because of his logical skill. The example doesn't settle the issue; but a rationale is needed for restricting *aretaic* evaluations from acts of skill.

³⁶ For a character-based account of responsibility for virtues, see Kupperman (1991). For how cognitive skills relate to responsible inquiry, cf. Polanyi (1958/2015) on "tacit knowledge."

a part of your character and expertise, explaining why you act by them. In chess, Tal's sacrifices are his partly because he is inventive; Botvinnik's positional game is his partly because he is prudent. And how each understands the position they see – even when playing each other – is informed by what possibilities that position affords: of sacrifice, for Tal, or of stepwise advance, for Botvinnik.³⁷

Responsibility for manifesting virtues doesn't presuppose reflective endorsement. As Hookway writes:

Virtues, we have said, enable us to respond to reasons: they provide a sensitivity to rational requirements in particular cases, and they are usually motivating. Although they can guide behavior through regulating deliberation and inquiry, their operation is not transparent to consciousness or open to reflective self-control. (Hookway 2000, p. 154)

While “reflective self-control” may exist, Hookway suggests it needn't³⁸ be operative when we manifest virtues or skills. If Hookway is right, to tie acts of virtue to reflective endorsement would put the cart before the horses, as virtues themselves may be “regulating deliberation.” Rather, what seems to be at play is a virtue of reflectiveness (Goldie 2008), which prompts reflection *only* when the situation calls for it.

³⁷ Perhaps these cases can be re-described; but possible examples suffice to make the conceptual point of why *aretaic* responsibility may be attributed to experts. I should note that I haven't seen *aretaic* evaluations for responsibility explicitly applied to experts so far in the literature. However, the possibility of such evaluations is guaranteed by the fact that, unlike virtue ethics, virtue epistemologies often treat virtues and skills jointly.

³⁸ Indeed, part of the reliability of one's acquired expertise presupposes it is tacit knowledge, operating seamlessly and for which the interference of reflection may be disruptive (Cianciolo et al. 2006). This would make little sense if responsibility for the manifestation of cognitive skills presupposed being able to reflectively endorse the resulting conceptions (also assuming that our cognitive architecture tracks what is most epistemically salient).

Suppose one thought we are nonetheless *able* to ask ourselves in reflection if the acts committed were appropriate even when the *occurrence* of reflection isn't called for. We may, in reply, point out that that ability to reflect plays no obvious role in explaining *why* the act committed was one of virtue if reflection isn't called for in order to commit the act.

Here, I don't develop a concept of responsibility for irreflective acts of virtue (in coming to understand included).³⁹ If such a concept failed to be coherent, that would be a larger problem for neo-Aristotelian virtue epistemologies generally than Carter and Pritchard's skeptical challenge to understanding. Rather, the point is only that (if the general challenge to responsibility for acts of virtue and skill fails) we may coherently deny (3), which requires reflective endorsement in order for one to qualify as being responsible for the formation of the conception by which one understands.⁴⁰ In the next section, I'll describe a set of cases best construed as counterexamples to (3) if the coherence of irreflective cognitive responsibility is granted.

³⁹ Nor do I propose a positive account of control, especially as it applies to one's skilled cognitive activity. For principled difficulties facing such accounts of control, see Fridland (2014).

⁴⁰ One might think, *pro* both Zagzebski and Carter and Pritchard, that we can only understand something when we are able to *explain* it, and that explanations require the ability to reflectively endorse one's conception of what is explained. However, just as in Section 5 we distinguished a reflective grip on what is understood from having a rational basis for holding one's conception, here we should distinguish the reasons one invokes in explaining what one understands from the verbal explanatory end-product. One may be sensitive to reasons in the way needed to satisfy the rational-basis requirement on understanding without thereby being required to be able to have a reflective grip on what one understands. Since what is at issue is virtue-based understanding, the point made here is essentially McDowell's (1979) point that one's virtuous sensibility to reasons in favor of acting as one does need not always (even possibly) be put into words. To say that verbal codifiability into rules for acting is *no constraint* on what gets to count as virtuous acting is less committal than Dreyfus' (2005) view that, say, chess-masters and seasoned drivers *don't* follow any rules.

7. An example

I deny (3) by providing examples in which one is cognitively responsible for understanding but doesn't reflectively endorse one's understanding. (The examples are some of the games played by Mikhail Tal, the Estonian 1960 world champion, more of which presently.) One doesn't always reflectively endorse one's understanding because reflective endorsement presupposes having a reflective grip on the rational basis for the conceptions by which one understands. Expert cognitive behavior (e.g., understanding chess positions) presupposes having a rational basis for conceptions one holds as an expert. However, expertise doesn't also require having a reflective grip over what one understands. So (3) is false; the chess examples will illustrate this.

Tal was one of the most tactical thinkers in 20th century chess. He was world champion for a year, 1960-61, until his predecessor Botvinnik won the rematch. I'm especially interested in Tal's early successes, 1957-1960, chronicled by Liepnieks (1961). Tal understood the positions he played, and was cognitively responsible in so playing. His inventiveness, surprising sacrifices, and an eye out for middle-game complications were well-known. Many of his spectacular moves withstand criticism. These facts indicate that he had a rational basis for the conceptions that constituted his understanding of the positions on the board. Reporting on his live performance, Tal says:

Calculation is only one side of it. In chess no less important is intuition, inspiration, or, if you prefer, mood. I for example cannot always explain why in one position this move is good, and in another bad. In my games I have sometimes found a combination intuitively simply feeling that it must be there. Yet I was not able to translate my thought processes

into normal human language.⁴¹

The games where Tal couldn't put into words the intuitive way he understood the positions on the board are evidence for thinking that Tal didn't have a reflective grip on those positions.⁴² Given the discussion in Section 6, such cases may plausibly be construed as counterexamples to premise (3) of Carter and Pritchard's skeptical argument.

We can use Tal's avowals to set up a telling instance of Carter and Pritchard's challenge. Consider *the availability heuristic* (Tversky and Kahneman 1973),⁴³ by which what comes to mind first is taken to be what is most relevant in context – its springing to mind first being an extra reason for thinking it is most relevant in context. Chess novices who don't stop and think before moving apply this heuristic, with amusing consequences. But experts often exemplify the heuristic as well.⁴⁴ They simply *see* the position for what it is, and move accordingly. They make what looks like a right move given their familiarity with positions of that kind.

Given the ubiquity of applying the availability heuristic – with varying results –, it is fair to say

⁴¹ Cited in Levinson (1994).

⁴² Tal wasn't inapt at explaining his own games; he chronicled his 1960 match with Botvinnik in a book now at its second edition. Rather, Tal seems to be saying that these explanations, while apt, are often retrospective rationalizations of what he there and then only *saw*, rather than articulated.

⁴³ Heuristics are shortcuts in problem solving, which lend themselves to ready use because they work most of the time. They often get the right result by cutting some corners. Part of the benefit, but also risk, that heuristics bring is the ability to bypass reflection altogether. This is why biases that apply heuristics inappropriately may threaten cognitive responsibility: even when the resulting conceptions are beneficial, their formation is arguably beyond one's control. It is a commonplace that some biases beset novices more; I am neutral on whether there are also biases that disproportionately affect experts.

⁴⁴ Chernev (1955, p. 429) records this quip: "Chess masters are wont to answer flippantly to the layman's question, 'How far ahead do you calculate, when making a combination?' They either say, like Réti, 'As a rule, not a single move.' Or, like Jaffe, 'I think one move ahead – but it is always the best move!'"

that although most of Tal's moves were due to his chess expertise (including appropriate uses of the availability heuristic), they could have easily been due to choosing a move because it happened to occur to him first (manifesting not his chess skill but, say, an unscrutinized overconfidence in his possession of that skill; a bias). After all, Tal was playing in largely the *same* way when his combinations and sacrifices were successful and when they were criticizable.⁴⁵ Through tongue-in-cheek remarks like “Some sacrifices are right; the rest are mine,” Tal seems to acknowledge this.

We can apply Carter and Pritchard’s challenge by saying that Tal couldn't have been responsible for understanding positions on the chessboard that could have easily had a *vicious* rather than virtuous epistemic etiology. Despite how spectacular his play was, a question should then linger as to whether Tal always genuinely understood those positions.

In reply to Carter and Pritchard’s challenge, successful and unsuccessful applications of the availability heuristic jointly support the view that not only biased thinking, but expert cognitive behavior too, is often intuitive and opaque to reflection. Intuitive expert performance such as understanding the position on a chessboard may be cognitively responsible even when one cannot endorse one's understanding in reflection.

⁴⁵ To put it in terms of current psychology of expertise (Gobet and Charness 2006), Tal had the same *knowledge base* each time (or even slightly better as time went by, given his youth), and he had the same (or even slightly better) *procedures to search* for the most advantageous feasible position. Notice, parenthetically, that neither knowledge base nor search procedures need to have been available to Tal during the game in declarative, propositional form – they could have been available as know-how, tacit knowledge manifested in his play rather than in reflective deliberation.

8. Conclusion

On virtue-based approaches to understanding (e.g., Zagzebski 2001), genuinely understanding something is primarily due to the epistemic agent's skills and virtues. Against such views, Carter and Pritchard (2016) raised the prospects of a skeptical position driven by the ubiquity and importance of biases in our epistemic lives: conceptions constituting how we understand something (e.g., the position on a chessboard) should, if arrived at primarily by intellectual virtues or skills, be accessible to reflection. This is because one is cognitively responsible for manifesting intellectual virtues and skills. But if the conceptions constituting our understanding could have easily been due to biases instead of virtues and skills, then a crucial aspect of their formation is reflectively opaque, and so we wouldn't count as fully responsible for understanding things the way we do. Whenever it could have easily been primarily due to biases instead of virtues and skills, our seeming understanding would not be genuine. As the threat of biases looms large for any degree of expertise, skepticism about understanding would be rampant.

I replied to Carter and Pritchard's skeptical challenge by denying a premise in the argument I construed them as advancing. Epistemic agents may be cognitively responsible for understanding something even when their understanding (e.g., of chess positions) isn't endorsed on reflection. This is typical, I have suggested, for experts: possibly acceding to reflection is not a requirement for responsible expert cognitive performance.

Chapter II

Mathematical understanding and “What if things had been different?” questions

According to Grimm (2014), we only understand a phenomenon if we know what other phenomena it depends on, and we identify dependencies according to how we answer “What if things had been different?” questions. I argue that this view meets with mathematical counterexamples. For, in mathematics, things couldn't have been different. I consider three replies Grimm may make, and argue they do not succeed.

1. Introduction

On Grimm's (2014) view, presented in Section 2, we only understand a phenomenon if we know what that phenomenon depends on. We understand how things stand because we know how they could have differed. That is, we know what depends on what because we're able to answer questions of the form “What if things had been different?”

However, in Section 3, I mention two instances of mathematical understanding that, I argue, are counterexamples to Grimm's view of understanding. The first is understanding the conclusion of a simple arithmetical argument on the basis of its premises. The second is understanding an axiom of set theory, the axiom of infinity. I argue that understanding this axiom, and understanding the conclusion based on the premises in the arithmetical argument given, are not mediated by our ability to answer “What if things had been different?” questions.

In Sections 4-6, respectively, I consider three replies Grimm may make, which interpret “What if things had been different?” questions differently: a purely modal approach, a grounding approach, and a modeling approach to answering such questions. I will argue that none of these approaches leads to re-interpreting the examples given in Section 3 so that they no longer undermine Grimm’s view. The conclusion of the chapter explores the wider epistemological consequences of this dialectical situation.

2. Understanding as knowledge of dependencies

In this section, I summarize Grimm’s view of understanding.⁴⁶ He writes:

Since understanding seems to arise from a grasp of all these different types of dependence, it might therefore be better [to] claim that understanding consists... of something like “knowledge of dependency relations.” (Grimm 2014, p. 341)

Consider an example, borrowed from Carr (1961). Suppose you set out to understand the 1905 Revolution (a social unrest brutally repressed by Czarist authorities in St. Petersburg in 1905). The occurrence of the 1905 Revolution depended on sundry factors: macro-economical, social, mass-psychological. Unions were prohibited and peasants newly arrived in the city had no workplace protections. Previous protests had been met with violence from riot police, and unrest was growing.

⁴⁶ This view is in the background of much contemporary theorizing of understanding, such as Zagzebski’s (2001) virtue epistemology of understanding, and Gopnik’s (1998) view of understanding as knowing the causal structure of the phenomenon understood. Closest to Grimm seems to be Greco’s view, for whom: “[T]o have understanding is to have systematic knowledge of dependence relations. To understand a thing is to be able to (knowledgeably) locate it in a system of appropriate dependence relations” (Greco 2014, p. 286).

Workers wanted acknowledgment of their problems. The newly created social stratum of industrial workers was disadvantaged in point of income and social status. Our understanding of the 1905 Revolution is partly constituted by our knowledge of the variety of its causes, and their interrelations. Efficient causation is only one kind of relevant dependence. Other kinds of dependence are temporal or logical relations (e.g., the 1905 Revolution preceded the 1917 Revolution; and it instantiates the concept of revolution). The 1905 Revolution *depends on* each of these because the Revolution might have not occurred at all, or might have occurred differently than it did, had any of these factors been absent or modified. Understanding is a matter of degree: we know more or fewer factors the Revolution depended on; and we know, in more or less depth, how each of them influenced the Revolution's unfolding.

The view that understanding requires knowing what the phenomenon understood⁴⁷ depends on flows from other of Grimm's commitments. Grimm (2010, p. 337) argues that “understanding is the goal of explanation,” so that at the end of explanatory inquiry the epistemic agent is able to explain why the phenomenon understood occurred. Moreover, Grimm (2014, p. 341) adheres to Kim's view that “explanations track dependence relations” (Kim 1994, p. 183).⁴⁸ And we identify what a phenomenon understood depends on by answering “What if things had been different?” questions:

⁴⁷ The word “phenomenon” is used broadly: any state of affairs, event, process unfolding, or property obtaining count as phenomena.

⁴⁸ As Kim elaborates: “my claim will be that dependence relations of various kinds serve as objective correlates of explanations. Dependence, as I will use the notion here, is a relation between individual states and events; however, it can also relate facts, properties, regularities between events, and even entities. We speak of the ‘causal dependence’ of one event or state on another; that is one type of dependence, obviously of central importance. Another dependence relation, orthogonal to causal dependence and equally central to our scheme of things, is mereological dependence: ...the properties of a whole, or the fact that the whole instantiates a certain property, may depend on the properties had by its parts” (Kim 1994, p. 183).

Following James Woodward (2003), one promising way to spell out the requisite notion of the kind of achievement that is needed – hence the kind of grasping that is involved in understanding—is in terms of having an ability to answer “what-if-things-had-been-different?” questions. To have an ability to answer questions of this sort, Woodward argues, is to be able to anticipate the sort of change that would result in the thing we want to explain (the *explanandum*) if the factors cited as explanatory (the *explanans*) were different in various ways. Brian Skyrms... expresses this same point in a more metaphorical way. According to Skyrms, grasping involves having an ability to anticipate that “wiggling” one variable will characteristically lead to a “wiggling” of another variable. (Grimm 2006, pp. 532-533)

For instance, if fewer poor people had moved from the countryside to the city; or if the government had instructed the riot police to react with caution to occasional conflicts on the streets and in factories; or if better workplace protections had been secured; or if the workers had been allowed to unionize and start negotiating their position (etc.), then perhaps the 1905 Revolution might have been less bloody than it was, or might have occurred later than it did.

We can put Grimm's view of understanding as follows:

- (1) If we understand a phenomenon, then we know how to explain why it occurs.
- (2) If we know how to explain why a target phenomenon occurs, then we know what other phenomena the target phenomenon depends on.
- (3) If we know what other phenomena the target phenomenon depends on, then we are able to answer “What if things had been different?” questions about the target phenomenon.

(4) So, if we understand a target phenomenon, then we are able to answer “What if things had been different?” questions about that phenomenon.

Or, as Grimm puts it, understanding involves “an apprehension of how things stand in modal space” (2014, p. 334). The argument is valid. Presumptive counterexamples may be produced to the thought that understanding requires knowledge or being able to explain what one understands.⁴⁹ But I am interested in what happens when everything goes well, epistemically speaking. So I will grant (1) for the sake of argument. Similarly, one may advance alternative conceptions of explanation, or resist any general view about scientific explanation; but I will grant (2) for the sake of argument.

I will argue in the next section that instances of mathematical understanding provide counterexamples to Grimm’s view⁵⁰ because, in mathematical cases, one *cannot* go about answering the question “Why is that the case?” by answering the question “What if things had been different?” For, in mathematical cases, things *couldn't* have been different. And considering patent impossibilities seems explanatorily idle rather than conducive to understanding.

As I granted (1) and (2) of Grimm's view, and since, in the next section, I offer two counterexamples to (4), I reject (3): the kind of understanding we have of phenomena we can explain doesn’t always require that we are able to answer “What if things had been different?” questions. In Sections 4-6, respectively, I consider three replies Grimm may make, intended to re-interpret the

⁴⁹ For discussion surrounding whether understanding a phenomenon requires being able to explain it, see Lipton (2009), Khalifa (2012) and Strevens (2013). The general conclusion to this dissertation suggests that how we should conceive the relationship between understanding and explanation largely depends on what notions of *both* understanding and explanation we work with.

⁵⁰ Grimm’s extending the “knowledge of causes” view to mathematical understanding contrasts with Woodward's (2003, p. 3) explicit disavowal of applying his view of scientific explanation to mathematical cases.

instances of mathematical understanding I mention so as to make them not conflict with Grimm's view. I will argue that the replies fail to achieve that end.

3. Two simple cases of mathematical understanding

In this section, I argue that two instances of mathematical understanding⁵¹ are counterexamples to Grimm's view.

Initially, Grimm (2006) had advanced a conception of understanding as “knowledge of causes,” according to which we only understand what we know the causes of. However, that conception involved efficient causation alone, and couldn't be generalized to cover mathematical understanding. To remedy that, Grimm (2014) advocates a theory of understanding based on dependencies in general rather than on efficient causation alone:

[G]iven how closely our notion of causation is tied to pushing and pulling – to exerting causal force – a perhaps more attractive strategy would be to demote the notion of causation from its central role and instead to appeal more generally to the notion of dependence (Grimm 2014, p. 341)

There are a number of problems with this passage that I bracket.⁵² However, I will argue that the resulting

⁵¹ The discussion to follow doesn't presuppose that mathematical understanding is a unified category of epistemic evaluation. There are a variety of mathematical projects worth pursuing, and they largely differ in their characteristic epistemic features. See Avigad (2008) and Hafner and Mancosu (2005) for different kinds of explanation in mathematics, and, presumably, different kinds of attendant understanding.

⁵² To consider only one issue: Our ordinary concept of causation *isn't* push-pull. We think demand may generate supply,

view faces counterexamples.

For a first counterexample, consider the understanding provided by the following argument:

(a) Every strictly positive natural number has a predecessor.

3 is a strictly positive natural number.

So, 3 has a predecessor.

The argument is valid. The premises are true, so is the conclusion. Moreover, the premises jointly *explain* the conclusion. One way of *understanding* the conclusion – the instance of understanding I'll discuss – is based on the premises. I will argue that Grimm's account of the understanding we have of the conclusion on the basis of the premises in (a) is false.

On Grimm's view, an answer to the question “Why are things so?” that provides an understanding of the state of affairs explained is obtained from an answer to the question “What if things have been different?” Given that the premises in (a) jointly provide an understanding of the conclusion in (a), that understanding has to be obtained *via* (b) below, since (b) answers the “What if things had been different?” question with respect to (a):

(b) If 3 had not been a strictly positive natural number, or if not every strictly positive natural number had a predecessor, then 3 might have not had a predecessor.

Statement (b) is a mathematical counterfactual. How to interpret mathematical counterfactuals is too large an issue to settle here.⁵³ It's clear that, when we understand why 3 has a predecessor on the basis of

Congress can keep the President in check, and a word can hurt our feelings; none of these are push-pull, and all are causal according to common sense.

⁵³ I submit it would be mistaken to think that all mathematical counterfactuals are trivially true given falsity of antecedents.

the facts that every strictly positive natural number has a predecessor, and 3 is such a number, what we are entertaining is (a), not (b). We have a grasp of the relevant dependencies (3 is a strictly positive natural number, hence shares their properties) without any appeal to counterfactuals like (b).

Here is a second example: the set-theoretic axiom of infinity, according to which there are infinitely many things. This claim is necessarily true if true at all. For there doesn't seem to be any room for asking what the existence of infinitely many things might be *contingent* on.⁵⁴ Suppose, for the sake of argument, that the axiom of infinity is true. It follows that no possible worlds exist where it is false. It's hard to see how the mere inexistence of possible worlds falsifying the axiom would *explain* the axiom or help us understand it.⁵⁵

I sketched two instances of mathematical understanding: the understanding afforded, *via* a simple arithmetical argument, of the conclusion by the premises; and the understanding we have of an axiom of set theory. These instances of understanding are importantly different. The first is within the grasp of grade school children; the second is more foundational. The first is arithmetical, the second is set-theoretical. The first is proof-based; the second is an axiom. If, as I contend, both are counterexamples

Consider: "If one [says] 'nothing sensible can be said about how things would be different if the axiom of choice were false,' it seems wrong...: if the axiom of choice were false, the cardinals wouldn't be linearly ordered, the Banach-Tarski theorem would fail and so forth" (Field 1989, pp. 237-238). Field is right. To appreciate whether the conditional is true, we have to grasp a bit of set theory and topology. That isn't trivial, regardless of whether the Axiom of Choice is true or false.

⁵⁴ In the text, I only need the claim that the axiom of infinity is necessarily true if true at all. There are traditional ways to deny this claim (such as Mill's suggestion that arithmetical truths are inductive generalizations generated by counting pebbles), but they seem to me to lack any plausibility.

⁵⁵ In keeping with (1), one may wonder in what sense we may explain the axiom of infinity when we understand it. Maddy (2011, pp. 113-137) addresses the issue, considering a form of explanation specific to foundational axioms, which she terms "elucidation." How elucidation may differ from other forms of understanding is beyond the scope of this text.

to Grimm's view, this shows that the problems with that view are quite general.

In Sections 4-6, respectively, I'll consider three replies Grimm might make, on which the instances of mathematical understanding I just mentioned aren't counterexamples to his view. I will argue that these replies don't succeed in explaining away the counterexamples. As Grimm's view of understanding was meant to apply to all cases of scientific understanding – mathematics included –, it should be abandoned because of the counterexamples it faces.

4. A modal approach to answering “What if things had been different?” questions

In this section, I argue against a reply that Grimm could make to the thought that the instances of mathematical understanding mentioned in Section 3 are counterexamples to his view. The reply is that we *do* have an answer to “What if things had been different?” questions. Namely, that they couldn't have been different. This answer is satisfactory with respect to the examples given in Section 3 because knowing how things *must* stand in the mathematical realm leads to an understanding of why things actually are that way.

I believe this reply misinterprets “What if things had been different?” questions. Consider a mundane example. In preparation for surgery, a patient is given an anesthetic. The dosage has to be right if the operation is to proceed seamlessly. We may ask: what if things had been different? Had the anesthetic dose been smaller, the patient might wake up on the operating table. Had the anesthetic dose been much larger, the patient could have a heart attack. By intervening on the anesthetic dose, we see effects in the patient's health condition. We know why the patient was sound asleep during surgery but

woke up fine afterwards: it was partly because he had been administered the right dose of anesthetic. Intervening on causes often alters their effects, so answering the question “What if things had been different?” in this context leads to discovering causal relations.

By varying the anesthetic dose, we vary its effects. In contrast, in mathematics, there’s nothing to vary, so long as we agree with Grimm’s (2014, p. 334) claim that knowledge of mathematical truths is *a priori*. So “What if things had been different?” questions clearly must play a different role in understanding matters mathematical than in understanding, e.g., how anesthetics work.⁵⁶

Once we admit that mathematical statements are necessarily true if true at all, then answering questions of the form “What if things had been different?” is trivialized. Compare “If every natural number has a successor, then $1+1$ exists” with “If infinite sets exist, then $1+1$ exists.” If both are necessarily true, then a purely modal approach will be unable to ascertain that the first expresses a genuine dependence whereas the second does not. Assuming that we do understand that 3 has a predecessor because all strictly positive natural numbers do, then either that understanding is not given *via* an explanation such as that in (a), *contra* (1), or that explanation doesn’t refer to dependencies between arithmetical facts, *contra* (2). If a purely modal approach to “What if things had been different?” questions in arithmetic trivializes their answers, given that banalities explain nothing, it is hard to see how a purely modal approach could underwrite our understanding of arithmetic.

Moreover, if answers to “What if things had been different?” questions in mathematics are

⁵⁶ We should distinguish Grimm’s view from a related one. “One way in which our wonderment about a phenomenon can be relieved is through a demonstration that it is necessary, that it could not be otherwise” (Glymour 1980, p. 32). Sometimes, we understand to actually obtain what we know obtains necessarily. This is surely true sometimes, and may apply often in matters mathematical. But it seems to obviate appeal to “What if things had been different?” rather than require it. Glymour is clear that this kind of understanding by knowing-it-is-necessary only holds *sometimes*.

trivialized, then finding them is no epistemic achievement. Even though we may readily produce such answers, it is no thanks to exercising a cognitive ability of ours. Grimm endorses (1), by which we only understand what we *know how* to explain. It follows that an element of cognitive achievement corresponds to coming to understand something new, and what is achieved is learning what the phenomenon understood depend on. This shows that, if (3) is assumed, (1) is undermined.

One might reply that counterfactuals like (b) can be interpreted so that they are non-trivial, by considering not only metaphysically possible worlds, but also metaphysically impossible worlds at which mathematical truths could possibly be false (Nolan 2014). Perhaps such a reply succeeds in making the attribution of a truth value to (b) non-trivial. But it doesn't alter the fact that (b) isn't required in order to understand the conclusion of (a) – that 3 has a predecessor – based on its premises: that all strictly positive natural numbers have predecessors, and 3 is such a number. And it is hard to see how an approach that considers metaphysically impossible worlds would *illuminate*, or enhance our understanding of, the axiom of infinity.

In the next two sections, I explore whether two other ways of interpreting “What if things had been different?” questions are conducive to explaining and understanding phenomena inquired into. I will argue that is not always so in the mathematical realm.

5. Grounding to the rescue?

Grimm mentions grounding (2014, p. 341),⁵⁷ suggesting the possibility of elaborating a reply to

⁵⁷ In fact, Grimm mentions Aristotelian formal causation as well. He regards the choice between appeal to grounding and

the claim that Grimm's view faces counterexamples relying on a grounding-based metaphysics. In this section, I argue that appeal to grounding is ill-suited to re-interpret the instances of mathematical understanding mentioned in Section 3.

The idea of appealing to grounding in order to fine-grain a modal approach to the dependencies grasped in understanding a phenomenon seems to be the following. To know what a phenomenon is caused by (*lato sensu*), or depends on, when we understand that phenomenon, is to know the grounds of that phenomenon. The idea that explaining why the phenomenon understood occurs is accounted for by answering the question "What if things had been different?" taken on a new guise: vary the grounds, and you will have thereby varied the grounded. By building a profile of how the grounds of a phenomenon vary, one can explain why the phenomenon occurs as it does.

Wilson argues that a general notion of grounding isn't theoretically helpful in illuminating what metaphysical dependencies are:

The problem here is not just that claims of Grounding (failure of Grounding) leave open some interesting questions; it is that such claims admit of such underdetermination — about whether the dependent goings-on exist, are reducible or rather distinct from the base goings-on, are efficacious, and so on—that even basic assessment of claims of

appeal to formal causation as alternative accounts of metaphysical dependence in general to be a terminological choice (a diagnosis I don't share). Grimm says that "Aristotle's notion of causation was more expansive, along the lines developed here" (2014, p. 341), referring to Aristotelian causes, in their formal, material, efficient, and final varieties. What formal causes are is quite obscure, but that problem can be sidestepped because these varieties of causes easily generalize to the plurality of dependencies. Mancosu (1996) chronicles how geometers gradually gave up talking about formal causes as they began articulating a notion of demonstrative proof anticipating those current nowadays. This is why I focus on grounding as a way to articulate Grimm's view of how understanding involves knowing what grounds what.

metaphysical dependence, or associated views, cannot proceed by reference to Grounding alone. As such, investigations into metaphysical dependence cannot avoid appealing to the specific ‘small-g’ grounding relations (again: type or token identity, functional realization, the classical mereological part-whole relation, the causal composition relation, the set membership relation, the proper subset relation, the determinable-determinate relation, etc.) that are capable of answering these crucially basic questions about the existential, ontological, metaphysical, and causal status of metaphysically dependent goings-on. (Wilson 2014, p. 540)

Wilson's point is the following. The theoretical construct of Grounding (capitalized to distinguish it from specific dependence relations) is supposed to capture metaphysical dependence in general. But this theoretical construct can't by itself help answer specific questions: “whether the dependent goings-on exist, are reducible or rather distinct from the base goings-on, are efficacious, and so on.” These questions, however, need answering. This is why one should look for another way to describe various metaphysical dependencies; and Wilson argues that part-whole, logical, causal relations etc. provide satisfactory answers to the questions that need answering. I endorse Wilson's criticism, and think it is the advocate of Grounding who has a burden of proof to show that Grounding can provide an illuminating account of metaphysical dependence.

How might a “small-g” grounding version of Grimm's view address the counterexamples in Section 3? Take the arithmetical counterexample first. The fact that every strictly positive natural number has a predecessor, and the fact that 3 is a strictly positive natural number, jointly *ground* the fact that 3 has a predecessor. This is an improvement over a purely modal approach, for now there are part-whole and logical relations that can be invoked as relations of metaphysical dependence going beyond whether

the mathematical facts in question hold with metaphysical necessity or not. Because 3 is *a part of* the strictly positive natural numbers, the fact that every strictly positive natural number has a property (having a predecessor) *logically implies* that 3 has that property too. These logical and part-whole relations can do justice to the understanding-giving and explanatory asymmetry of the premises and conclusion of (a), improving over a purely modal elaboration of Grimm's view.

However, a grounding elaboration of Grimm's view has trouble accounting for (b). For (b) is an answer to “What if things had been different?” Whereas grounding relates to metaphysical necessity, and it *doesn't* seem to make sense to ask what it might have been for metaphysically necessary facts (that every strictly positive natural number has a predecessor, that 3 is such a number, etc.) to be different from what they are. We lack an answer to the question of what it is to *vary* mathematical grounds. By the same token, we can't vary facts (taken to be metaphysically necessary) such as the existence of infinitely many things.⁵⁸ So a grounding-based reply in support of Grimm's view doesn't succeed in addressing the understanding we have of the axiom of infinity either.

I have considered an interpretation of “What if things had been different?” questions on which answering them indicates what grounds what, for varying the grounds would vary the grounded. And I have argued that, while this improves over a purely modal approach to “What if things had been different?” questions, it doesn't rescue Grimm's view from counterexamples.

⁵⁸ One might propose that the existence of infinitely many things is grounded in the continuity of spacetime. However, that spacetime is continuous (if it is) seems to be a claim of less modal strength than what the axiom of infinity requires, for it is a claim in the foundations of physics, not of set theory.

6. Different models

Grimm (2014, p. 338) argues that “mental models” may play a role in the cognitive activity that underwrites understanding. This raises the question of whether a modeling approach to understanding⁵⁹ could rescue Grimm’s view from counterexamples. The guiding thought would be that, when we ask how things mathematical could have differed, we contemplate alternative models we use. My reply is that, while we may at times contemplate a variety of models, that isn’t always required for understanding the phenomena theorized.

This text doesn’t discuss a modeling approach to understanding *per se*, but only inasmuch as it offers a new interpretation to “What if things had been different?” questions asked in order to identify what a phenomenon depends on, so as to explain and understand that phenomenon.⁶⁰ A modeling

⁵⁹ Models, here, are conceived broadly, so as to include both representations internal to a thinker’s mind – her memory of concepts or schemata – and external representations, models understood as cultural artifacts to be found in textbooks or laboratories. I depart from Grimm’s reference to “mental models” in order to consider a view as close to Grimm’s as possible that yet remains non-committal on the cognitive activity that underwrites understanding, for such psychological speculation would raise questions of plausibility on its own.

⁶⁰ We should distinguish the view considered in this section from Saatsi’s (2018) view. Saatsi writes that “genuine explanations are underwritten by explanatory dependencies in the world. This is the basic factivity requirement of the counterfactual-dependence account. Explanatory understanding, in turn, can be construed as an agent’s ability to make correct counterfactual what-if inferences. [And] what matters for explanatory progress is that understanding-providing theories and models *de facto* latch onto reality in appropriate ways, so as to satisfy explanations’ basic factivity requirement.” (Saatsi 2018, pp. 14, 2). What Saatsi seems to be suggesting is that we explain phenomena by drawing inferences from the models we have of those phenomena. And models represent, not just what there is, but also how the target systems would behave in different conditions. This differs from the view attributed to Grimm here, on which “What if things had been different?” questions consider, not varying phenomena given a fixed model, but varying models of presumptively the same target phenomenon.

approach wouldn't eliminate the role "What if things had been different?" questions play. Such questions are important in the natural and social sciences, where counterfactual scenarios are a good guide to testing. Rather, we answer "What if things had been different?" questions *differently*, by considering how our models could differ in how they represent the same target phenomenon, even when that phenomenon obtains necessarily.⁶¹ The word "model" has an astonishing variety of uses.⁶² However, in order to speak to the counterexamples in Section 3, the view has to presumably consider models in the model-theoretical sense.

In the case of arithmetical understanding mentioned in Section 3, answering "What if things had been different?" questions has to meet two tasks. First, that we can make sense of the mathematical counterfactual (b). Second, that it is somehow *via* interpreting this counterfactual that we gain the understanding in (a) of the fact that 3 has a predecessor. I believe that the first task is met but that the second is not.

Here are two models relative to which we may make sense of (b), or non-trivially assign it a truth

⁶¹ On a modeling approach to "What if things had been different?" questions, it doesn't matter that the antecedent of (b) may not pick out any metaphysically possible worlds. For we can ask: "What if things had been different – logically, set-theoretically, arithmetically, etc.?" And these questions would be neutral on matters metaphysical, showing that the third way of interpreting "What if things had been different?" questions is superior to the first two ways of interpreting them, sketched in Sections 4-5.

⁶² After giving several quotes from science textbooks illustrating different uses of the word "model," Suppes writes: "The first of these quotations is taken from a book on mathematical logic, the next two from books on physics, the following three from works on the social sciences, and the last one from an article on mathematical statistics... One of the more prominent senses of the word missing in the above quotations is the very common use in physics and engineering of "model" to mean an actual physical model as, for example, in the phrases "model airplane" and "model ship" ... It would, I think, be too much to claim that the word "model" is being used in exactly the same sense in all [these contexts]." (Suppes 1960, p. 3) In arguing that "the meaning of the concept of model is the same in mathematics and the empirical sciences" (p. 4), Suppes is aware that he is going against first appearances.

value. First, consider a model of first-order predicate logic that has only two elements. This model would *not* also be a model of first-order Peano arithmetic. In this model, (b) is true because its antecedent is false. But the justification is not due to any metaphysical impossibility. Rather, the justification has to mention the *model* in which we interpret (b). Second, consider a model of two elements, but this time as a model of the positive integral domain \mathbb{Z}_1 . On that model, for instance, $(4)_{\mathbb{Z}_1} = (2)_{\mathbb{Z}_1} = (0)_{\mathbb{Z}_1}$ and $(5)_{\mathbb{Z}_1} = (3)_{\mathbb{Z}_1} = (1)_{\mathbb{Z}_1}$. Mimicking the successor function for natural numbers, in \mathbb{Z}_1 the successor of 0 will be 1, the successor of 1 will be 0, and so on. “3” is interpreted as $(3)_{\mathbb{Z}_1}$, and we can ascertain that it does have a predecessor. Again, (b) will come out true because its antecedent is false. But, again, the falsity of the antecedent of (b) and the truth of (b) are established partly with reference to the model we interpret (b) in.

I’ll now argue that showing (b) is coherent and can be interpreted non-trivially doesn’t *also* show that the understanding (a) delivers of why 3 has a predecessor has to be mediated by (b).

We might think as follows. We can test the validity of the argument in (a) by considering models, trying to falsify the conclusion and see if the premises are thereby falsified or not. Since the argument in (a) is valid, no such model will be forthcoming. And it is open to say, with Grimm, that we might better *appreciate* the validity of the argument in (a) by considering models. In considering such models, *volens nolens* we consider models that would give (b) the same truth value, albeit for different reasons. Whether we put the matter in terms of the verbalized subjunctive conditional (b) is irrelevant. What matters is that we entertain such models in testing if (a) is valid.

There is some plausibility to this line of thought, but it surely cannot hold in all cases, and we have no special reason for thinking it holds in the case of (a)-(b). The key point to notice is that we may have a purely syntactic understanding of why 3 has a predecessor. We grasp that all strictly positive

natural numbers have predecessors, and that 3 is a strictly positive natural number, we grasp *modus ponens* and universal instantiation, and can thereby deduce the conclusion in (a), that 3 has a predecessor. We need appeal to no models in doing so.

The reply generalizes: we may have a purely *proof-theoretic* understanding of a conclusion without appealing to modeling. To be sure, valid arguments are such that, in any model where the premises are true, so is the conclusion. But mention of models is indirect, and it is unclear why understanding the conclusion of a proof *could not* be achieved without reference to models.

I am not advocating a proof-theoretical approach to mathematical understanding. Intuitively, the proof-theoretical approach and the model-theoretical approach are complementary, and precisely which of them, and in what contexts, better serve understanding should not be prejudged. But this is precisely what (3) does in implying that we can only understand the conclusion in (a) *via* (b).

In this section, I argued that a modeling approach to “What if things had been different?” cannot rescue Grimm’s view from counterexamples.⁶³

⁶³ To simplify things, the text focuses on the first counterexample to Grimm’s view in Section 3. As for the second counterexample in Section 3, we might think that two images support the axiom of infinity. Either infinitely many points exist in spacetime (so infinitely many concrete individuals). Or we may built the iterative hierarchy of sets, which is as close as we may come to intuitiveness without running into inconsistency. (Field 1980 and Boolos 1971 are the classical sources for these two images of sets.)

Asking how things might have been different would, then, amount to asking whether we may contemplate models that retain plausibility despite departing from these two images of what sets are. Finitism, for instance (even if *per impossibile*), is a live theoretical option. The point is, not to decide any issues with respect to infinity, only to point to how basic theoretical and modeling choices inform how we understand the axiom of infinity.

In reply, notice that calling the claim that there are infinitely many spacetime points, or the iterative theory of sets “images” plays no essential role. We might as well *deduce* the axiom of infinity from such claims and theories. Suppose we wrote the model theory for the Zermelo-Fränkel set theory in a meta-language of our choice, so that the “intuitive”

7. Conclusion

On Grimm's (2014) view, we only understand a phenomenon if we know what its occurrence depends on. On this view, we know what depends on what, and are enabled to explain target phenomena, by answering "What if things had been different?" questions. However, I developed two counterexamples to Grimm's view of understanding, and argued against three replies Grimm may make.

This dialectical situation is important because Grimm's view considers understanding in full generality, whereas the counterexamples I mentioned are instances of mathematical understanding. The fact that Grimm's view meets with these counterexamples may suggest that different instances of understanding may vary significantly in their epistemic profile according to *what* is being understood: bits of mathematics, the occurrence of revolutions, or how the dosage of an anesthetic contributes to the success of ensuing surgery. The variety of possible interpretations of "What if things had been different?" questions, and how we may answer them, reflects this. Whether an informative general, topic-neutral, epistemology of understanding does justice to the epistemic particularities of different instances of understanding is still an open question.

iterative conception of sets is simply a theory in one's chosen meta-language, all of whose models are also models of the object-language Zermelo-Fr nkel set theory being interpreted. All this is to say that no obvious benefit follows from appeal to "images" as opposed to just more theory.

Chapter III

Experience and understanding

How do conscious experience and understanding relate? I start from two recent views, those of Bourget (2017) and Lynch (2017). Relying on objections to both, I propose my own view, according to which any instance of understanding occurs in one of three modes. Each mode may be experienced: Moments of coming to understand are conscious in insights, or “Aha!” experiences. Exploiting an understanding one already possesses in a discursively articulated way is conscious in intuitions. And exploiting an understanding one already possesses but which one cannot discursively articulate (an otherwise tacit understanding) is made manifest in experiences of fluency.

1. Introduction

People often report, accurately, that they understand a mathematical theorem, or they can see why it rained heavily after protracted drought, or that they can only now make sense of what a grimace meant. When we ask what understanding is, part of what we wish to know is what makes such reports true.

The term “understanding” is used in many ways: linguistic, perceptual, practical, emotional, and intellectual. Let's illustrate these uses: We know what one means when we understand what they say. When we understand what we see, we recognize objects in our visual scene according to the concepts we have. We feel empathy for someone when we understand their plight. We are finally able to dance well

once we understand how to dance. And we intellectually understand something we make good enough sense of, that is intelligible and coherent enough for us. In what follows, I only discuss intellectual understanding.⁶⁴

In this text, I examine how conscious experience and understanding are related. We may consider whether all understanding occurs consciously; and we may consider whether there is something it is like to understand regardless of what is understood. In what follows, I discuss two recent views advanced by Bourget (2017) and Lynch (2017) concerning these two questions, respectively. In Sections 3-5, I discuss Bourget's view. And in Sections 6-8, I discuss Lynch's view. I argue that each view makes a questionable assumption. Then, in Section 9, I present my own view of how conscious experience and understanding relate. I argue that my view does not make the questionable assumptions Bourget's view and Lynch's view make, offering a *pro tanto* reason to endorse my view over theirs.

2. Setting the stage

Before it can be properly answered, the question of how understanding and conscious experience are related needs sharpening. I follow Bourget (2017) in drawing two distinctions,⁶⁵ that between grasp

⁶⁴ I leave it open how often other uses of the term "understanding" properly apply to instances of intellectual understanding. For instance, religious understanding presumably involves practical, emotional, and intellectual aspects, which would make it appropriate to apply the term "understanding" in several of its respective uses.

⁶⁵ "Experience," in this text, means conscious experience. Experiencing understanding is not a higher-order experience about our first-order understanding. Rather, experiences of understanding are first-order experiences with a certain content, and that content is what we understand.

and understanding, and that between occurrent and dispositional mental states.

Bourget writes that:

[I]t seems that grasping is independent of such epistemic factors as justification and truth. We can, at any case, distinguish a purely mental component to the state [of understanding]. That is what I mean by “grasping.” (Bourget 2017, p. 292)

We may apply the method of doubt in order to distinguish grasp from understanding. For instance, suppose the conception by which you understand planetary motion in our galaxy partly consists in your belief that Kepler’s law is approximately true, and the trajectories of planets approximate ellipses. Now suppose an all-powerful evil demon deceived you, so that the laws of nature didn’t actually include that planetary trajectories approximate ellipses, though it may seem so to you. Your conception of planetary motion would then be false, and unwarranted, yet it may nonetheless underwrite a grasp of planetary motion on your part. For, truth and warrant aside, the “purely mental component” of your understanding of planetary motion would be preserved in an evil-demon scenario.⁶⁶

The grasp/ understanding distinction matters because only mental states or episodes may be conscious. It follows that, if understanding relates to consciousness in any important way, it has to be *via* grasp.

Except in special circumstances, we don’t forget or instantly lose the grasp we have of traffic rules, political institutions, or basic human anatomy. Understanding and grasp are long-term mental

⁶⁶ What is internal to a mind – or “purely mental” – is controversial. For several ways of drawing the internal/ external distinction, see Gertler (2012). I think Bourget only assumes that some such distinction is conceptually coherent, without committing to any specific theory of how to draw that distinction.

states. When and how does what we grasp come to mind?

To address this question, Bourget draws a second distinction, between an occurrent grasp and a dispositional grasp, analogous to that between dispositional and occurrent belief.⁶⁷ Dispositional belief endures. By contrast, occurrent belief is a mental episode. If an occurrent belief shares content with a dispositional belief of the same thinker, the former typically manifests the latter.

Suppose you have a dispositional grasp of how economic cycles develop. If queried, you could say what you remember from Keynes' work by way of explaining how boom and bust alternate, and how the greater the difference between real value and market value during a boom, the more damaging the bust to follow often is. Naturally, you may have this dispositional grasp of economic cycles even when you're not thinking about it.

Bourget conjectures that dispositional grasp is nothing other than the disposition to occurrently grasp. When called for, one may bring to mind one's understanding of something – say, of economic cycles, in order to begin explaining the 1929 stock market crash. Bourget writes:

More generally, it seems that non-occurrently grasping *P* is simply a matter of being disposed to occurrently grasp *P* as required. This view, which mirrors the above view of non-occurrent belief, gives explanatory priority to occurrent grasping over non-occurrent grasping. (Bourget 2017, p. 292)

⁶⁷ “Consider, for example, the proposition that cats have whiskers. Now that I have mentioned that proposition, your belief in it is occurrent. But you have believed it continuously for decades in the non-occurrent sense” (Bartlett 2018, p. 1). By what criteria to draw the occurrent/ dispositional distinction is quite controversial. For discussion, see Bartlett (2018). Again, I only rely on the assumption that some such distinction is conceptually coherent, without assuming any specific theory of it.

Even if non-occurrent grasp of P is the disposition to occurrently grasp P , grasp may often not occur fully all at once. Bourget illustrates this (p. 306) with a seminar example: in the course of a conversation about physicalism, asking someone what they mean by supervenience, and receiving a satisfactory answer, may show that they knew what they meant all along, or that they had a firm grasp of the notion albeit their earlier remarks had not intimated that.⁶⁸

Only occurrent mental states or episodes can be phenomenally conscious. So, in order to determine how conscious experience and dispositional grasp are related, we need to determine how conscious experience relates to occurrent grasp.

With the grasp/ understanding and occurrent/ dispositional distinctions in mind, we can replace our original question with a more precise one. We had asked how understanding and conscious experience relate. We may now ask: How do occurrent grasp and conscious experience relate?

3. Grasping by experiencing

In this section, I present Bourget's view of grasp and understanding, and how they are related to

⁶⁸ The full quote is this: "Most of the time, we do not fully and occurrently grasp the complete and precise contents that we express or may plausibly be attributed. Imagine for example that you are in the middle of a philosophical conversation about the alleged supervenience of the mental on the physical. Your interlocutor suddenly asks what you mean by 'the mental supervenes on the physical.' After a pause, you might find yourself replying something along the lines of 'I mean *necessarily, any two objects that have the same physical properties have the same mental properties.*' When you 'unpack' a thought in this way, it seems that you form a better occurrent grasp of its content than you had a moment before. This process can go through many iterations. For example, your interlocutor might ask what you mean by 'object' or 'physical,' which is going to require some thought" (Bourget 2017, p. 306).

conscious experience. Bourget's central claim is:

To occurrently grasp *P* is to have a phenomenal experience with *P* as content. (2017, p. 303).

Bourget glosses this by saying that “to have a thought with content *P* while grasping *P just is* to experience *P*” (*ibidem*).⁶⁹ So, consciously experiencing *P* is (said to be) a necessary condition for the occurrent grasp of *P*.

This view clarifies what the relation between understanding and conscious experience is. The “purely mental component” of dispositional understanding is dispositional grasp. Dispositional grasp is acquired and subsequently manifested in episodes of occurrent grasp.⁷⁰ And to occurrently grasp *P* is

⁶⁹ If Bourget means to say that occurrent grasp is *identical* to a phenomenal experience whose content is the grasped content, this seems to commit Bourget to the existence of constitutively cognitive phenomenology *regardless* of the sensory support for the concepts deployed in the thoughts by which we grasp, which Carruthers and Veillet (2011) have argued isn't possible. In contrast, it would initially seem that the phenomenal theory of grasp is orthogonal on debates concerning constitutively cognitive phenomenal contents. The phenomenal theory of grasp should, then, be able to be formulated in a neutral way on that issue.

In reply to this line of thought, Bourget writes: “According to the phenomenal theory, grasping a proposition is a matter of having a phenomenal experience that has the proposition as its content. This tells us what grasping is, but this leaves open what is involved in *thinking* something with grasping or understanding, which was our original question. One possible view (arguably the most natural) is that to have a thought with content *P* while grasping *P just is* to experience *P*. Another view is that grasped thoughts involve experiences with relevant contents without being identical to them. The choice between these two views seems to turn more on how one wants to use the word ‘thought’ than on what is involved in grasping: unless one has reasons for not wanting to call experiences that occur in cognition ‘thoughts,’ I don't see why one should not endorse the first view. In any case, I will focus on the first and simpler view here.” (Bourget 2017, p. 303) Like Bourget, I hesitate whether what is at stake is how to use the words “thought, grasp” or both. It seems, however, that, even on the weaker version, consciously experiencing *P* is necessary for the grasp of *P* to occur.

⁷⁰ In fact, things are less simple, but the simplification is harmless for the purpose of developing my objection to Bourget's view. Two different situations may obtain in grasping something, by Bourget's lights. We might grasp in a *fulguration*, so to speak. For instance, I may understand what a grimace meant in the context of a tensed conversation, but be absorbed

necessarily associated with consciously experiencing *P*. So, conscious experience is necessarily associated with the manifestation of the “purely mental” dispositional “component” of understanding.

Bourget's view also does justice to the commonsense thought that experiences of understanding vary significantly. Suppose, that, in conscious thought, one has a grasp both of the genesis of Lutheranism and of bleak projections for climate change. What it is like to think of Lutheranism shares little with what it is like to think of climate change. To the extent that the phenomenal characters of concepts involved in grasping Lutheranism and climate change respectively differ, it follows from Bourget's view that what it is to occurrently grasp each of these will also differ. Part of the variety in experiences associated with understanding, Bourget's view suggests, is due to differences in what is understood.

Surveying a range of cases, Bourget concludes that they make most sense, or are best explained, by agreeing that “[t]here is at least a rough, *prima facie* correlation between the availability of relevant episodes of phenomenal consciousness and the capacity to grasp a given content.” I think this is a modest and fair summary of what evidence we may invoke. However, a “rough, *prima facie* correlation” between conscious experience and occurrent grasp falls short of showing that consciously experiencing *P* is a necessary condition for occurrently grasping *P*. This elbow room will be important in assessing my criticism of Bourget's view in the next section.⁷¹

by how the conversation is going and immediately forget that specific grimace. My grasp of the grimace would occur at that moment without leading me to acquire a stable and robust disposition to interpret my interlocutor's mimic.

Another situation may also obtain. Bourget says: “Most of the time, we do not fully and occurrently grasp the complete and precise contents that we express or may plausibly be attributed” (p. 306). So it is possible, for most contents we dispositionally grasp, that our dispositional may *never* be on full display in any single instance of its manifestation. Think, for instance, of the foremost commentator of a poet, who grasps that poet's universe of motifs and metaphors as no one else, yet who in no commentary succeeds in imparting one's full grasp.

⁷¹ Bourget notes the difference: “The cases discussed... exemplify at least a rough correlation between the ability to grasp a

4. Is unconscious grasp as much as possible?

In this section, I object to Bourget's view. The objection, in a nutshell, is that we have no good reason to think that unconscious occurrent grasp is impossible. In Section 9, I will offer my own view of how understanding and conscious experience are related. Unlike Bourget's view, my view doesn't presuppose that unconscious occurrent grasp is impossible – nor do I deny it. If the assumption at issue is questionable, we have a *pro tanto* reason to endorse my view over Bourget's view.

Of course, if *all* occurrent mental states were conscious (and perhaps necessarily so), occurrent grasp would be conscious as well. Although several arguments have been proposed in favor of views on which all occurrent mental states are conscious, it would be incumbent upon resulting views to explain away what have often been thought to be counterexamples: cases of unconscious perception (in experimental psychology), unconscious desires (in psychoanalysis) and unconscious beliefs (in the psychology of reasoning and of bias). In what follows I bracket the issue because it doesn't seem to pertain to grasp and understanding explicitly. The question I pursue is why, assuming that other kinds of mental states may possibly occur unconsciously as well, we would have any *special* reason to believe that unconscious occurrent grasp is impossible.

Consider a case Bourget himself mentions:

The opaque proof. You are trying to understand a proof. You know (because you have

certain content and the ability to have related phenomenal experiences. This *suggests* a view according to which grasping constitutively involves phenomenal consciousness” (Bourget 2017, p. 296, my italics).

been told by someone you trust) that the conclusion follows from the premises. You have gone through all the steps of the derivation. Still, you don't really *see* how the conclusion follows from the premises. At last, after going through the steps numerous times, you finally see it. Here the proposition grasped is along the lines of <such and such follows from such and such>. (Bourget 2017, p. 288)

Bourget avows that “[t]he precise role of experience is less clear in the proof case” (p. 294), and that seems entirely correct. There may often be “a phenomenological change when the parts snap together” (*ibidem*). But the question is whether this *must* happen in order for grasp of what follows from what to occur. Otherwise, the correlation between conscious experience and occurrent grasp would only be happenstance; whereas, on Bourget's view, “to have a thought with content *P* while grasping *P* *just is* to experience *P*” (p. 303).

If we see logic students competent in producing correct proofs, apt in evaluating other proofs, and typically doing so after a few moments' thought, doubtless the students undergo many conscious experiences meanwhile, but *a priori* there is little reason to think that they *must* have experiences with the content that so-and-so follows from such-and-such, unless they themselves volunteer introspective reports stating that is the case. Unconscious occurrent grasp is conceivable, and that gives us a *pro tanto* reason to think it is also possible.

In reply, Bourget notes that “it seems likely that most of our cognitive experiences must go unnoticed” (p. 307) and that “most of our grasped contents are only non-occurently grasped” (p. 306), and I grant that these remarks are plausible. But the onus of proof for Bourget is to show that it isn't *possible* for unconscious grasp to occur, since that is what follows from claiming that consciously

experiencing P is a necessary condition for the grasp of P to occur. And that hasn't been shown.⁷²

In the next section, I consider (and object to) a positive argument that Bourget offers in favor of the view that grasp can only occur consciously.

5. Inference, grasp, and conscious experience

In this section, I present and criticize Bourget's argument for the claim that occurrent grasp requires conscious experience of what is grasped. What seems to be Bourget's primary argument for the phenomenal theory of grasp is that it does better justice to grasp than an alternative he dubs "the inferential view of grasp," according to which:

An occurrent thought t with content P is an occurrent grasping of P to the extent that t is appropriately inferentially connected to other mental states of the subject. (Bourget

⁷² One may plausibly think that Bourget may support his view that occurrent grasp is necessarily conscious by appeal to the phenomenal intentionality research program (Kriegel 2013). According to that program, if there are any unconscious occurrent mental states, their intentionality is derived from that of conscious mental states.

Bourget's view is certainly consistent with that program: "If consciousness is the engine of reason and grasping a proposition is a matter of bringing it within the purview of consciousness, we should expect grasping to play a central role in the generation of rational behavior" (Bourget 2017, p. 312).

This quote is, however, quite controversial. First, it is unclear why, if the support for Bourget's view comes *via* consciousness being "the engine of reason," that isn't indirect support for the view Bourget argues *against*, viz. the inferential theory of grasp. Second, it is hard to see how occurrent grasp could "play a central role in the generation of rational behavior" if Bourget were also right to claim that "grasping seems to be a purely mental phenomenon that is independent of one's epistemic state" (Bourget 2017, p. 291). If stripped of epistemic properties, it is hard to see *in virtue of what* grasp would play a central role in rationality.

2017, p. 298)

Bourget raises a host of objections (pp. 298-299) to the inferential view⁷³ that I am ready to concede. But I believe that Bourget may be setting up a false dilemma, and both the inferential view and the phenomenal view of grasp may ultimately be unsuccessful.

It seems plausible to think that, if we are to characterize grasp, we may do so *either* in terms of what it is like to have experiences whose content is grasped, *or* in terms of what epistemic properties grasping has, so that we should be able to distinguish between, say, merely believing *P* and believing *P* with a full grasp of its content. (I don't think the phenomenal-epistemic disjunction is exclusive.)⁷⁴

However, it is far from clear that presumptive epistemic properties of grasp can only be inferential. Here is how Sosa and Zagzebski describe understanding:

Beyond 'animal knowledge' there is a better knowledge. This reflective knowledge does require broad coherence, including one's ability to place one's first-level knowledge in epistemic perspective. But why aspire to any such thing? What is so desirable, epistemically, about broad coherence? Broad coherence is desirable because it yields integrated understanding[.] (Sosa 1997, p. 422)

Understanding is a cognitive state that arises from *technê* ... The person who has

⁷³ Bourget rightly attributes this view to Stephen Grimm; see, e.g., Grimm (2016). He might equally well attribute it to Lynch (2017); see below for a discussion of his view.

⁷⁴ It might, perhaps, be more plausible to consider how intuitions (as instances of conscious grasp) and inference *relate*, rather than exclude each other. Thus, A. C. Ewing writes that "a philosophy may, like Spinoza's, find its culminating point in an intuition which goes beyond the reasoning that has preceded it. Though it might never have been possible but for that reasoning" (Ewing 1941, p. 36). It would seem that, on Ewing's view, both inference and intuition may sometimes contribute to understanding.

mastered a *technê* understands the nature of the product of the *technê* and is able to explain it. (Zagzebski 2001, p. 240)

Let's suppose that Zagzebski and Sosa discuss the same mental state of understanding and attribute different epistemic properties to it. Here are four such epistemic properties grasp *might* have, according to the two:

- (i) accessibility of grasp to reflection;
- (ii) doxastic coherence of what is grasped with our web of beliefs;
- (iii) the ability to explain what is grasped;
- (iv) the manifestation of intellectual virtues and skills in the formation of grasp.

It is unclear that any of the epistemic properties mentioned in (i)-(iv) qualify as “suitably... inference-like dispositions,” as Bourget (p. 295) would have it. If these epistemic properties are not best construed in inferential terms, then requiring them for grasp would lead to a view of grasp distinct from *both* Bourget's phenomenal view and the inferential view he opposes.⁷⁵

One reply open to Bourget is to say that (i)-(iv) mention epistemic properties of understanding but *not* of grasp. In drawing the understanding/ grasp distinction, Bourget noted that, since grasp is the “purely mental component” of grasp, it follows that grasp is independent of the truth of what is grasped (though perhaps genuine understanding is not). *If* one has an externalist conception of justification, that would also be a property of understanding but not of grasp. (Grant that assumption for the purpose of

⁷⁵ Alternatively, if Bourget means to suggest that all epistemic properties in (i)-(iv) should somehow be accounted for in terms of *inferring*, an account of inference is due, and will likely be controversial. For some of the pitfalls besetting offering any account of inference, see Boghossian (2014).

argument.) However, *contra* Bourget, it doesn't follow that "grasping seems to be a purely mental phenomenon that is independent of one's epistemic state" (p. 292). For the epistemic properties of understanding mentioned in (i)-(iv) seem "purely mental" as well, and would hence carry over as epistemic properties of grasp.

Other than giving examples (vivid and persuasive, but falling short of showing that unconscious occurrent grasp is *impossible*), criticizing the inferential view of grasp as though it were the only epistemic view of understanding worth its salt⁷⁶ was, as far as I could tell, the only positive argument Bourget offered in his (2017) for the claim that grasp can only occur consciously.

It seems that Bourget has provided no good reason to rule out the possibility that, in some circumstances, grasp and understanding *may* occur unconsciously. As Strawson (1994, p. 10) puts it, perhaps "talk of understanding-experience may not be appropriate in all cases in which it is correct to say that someone has understood something."

I don't *deny* Bourget's main claim that occurrent grasp may only be conscious. Rather, in Section 9, in presenting my own view of how conscious experience and grasp relate, I stay non-committal on whether grasp may only occur consciously. Not taking sides on a claim so controversial seems more prudent, and is a *pro tanto* reason to favor my view over Bourget's.

In Sections 6-8, I turn to an alternative, Lynch's view of how understanding and conscious experience relate. I will argue that Lynch's view makes a different questionable assumption. Then I will propose my own view in Section 9, and argue that it avoids making the controversial assumptions

⁷⁶ Bourget considers (at pp. 296-297) and objects to views of grasp other than the phenomenal and the inferential view. However, he does not seem to consider the family of views I point to in the text.

Bourget's view and Lynch's view make.

6. Insight and understanding

In this section, I present Lynch's (2017) view about the role of insights in understanding. Lynch gives a robust epistemic description of what states of mind qualify as states of understanding:

A state of some agent plays the understanding-role with regard to some subject when its content concerns dependency relations between propositions or states of affairs relevant to the subject; it is conducive of the agent's ability to offer justified explanations of the relevant subject; and it disposes the agent to make further justified inferences both factual and counterfactual about the subject.

Such a state would presumably have positive epistemic status and probative force.

(Lynch 2017, p. 199)

So conceived, how does understanding relate to conscious experience? Lynch writes:

My hypothesis is that an analysis of the cognitive act of coming to understand can help shed light on why we are tempted to say that understanding (the state) is a cognitive achievement, and why it involves an activity like grasping. (p. 201)

For Lynch, "the causally prior act of coming to understand...is partly definitive of the state it produces" (p. 200). How, then, do we come to understand?

[T]he moment of coming to understand can involve sudden insight. Such moments are often called "aha moments"... Of course, most acts of understanding do not require the

sudden novel inspiration.... But all of them do involve some level of insight. Having such an insight is part of why understanding is fundamentally a creative act. (p. 202)

Why should coming to understand always involve “some level of insight”?

Coming to understand has a particular phenomenological appearance. In cases of sudden insight, this phenomenological aspect of creativity either constitutes, or leads to, the “eureka” feeling. But creative acts can be surprising even if they do not necessarily provoke that “aha” moment.... Even when coming to understand happens gradually over time it still feels “new” – as if you couldn’t have understood it prior to that moment. It feels as if you’ve made forward progress. (p. 203)

Here is how understanding and conscious experience are related, on Lynch’s view. The act of coming to understand is “partly definitive” of the understanding that ensues, and “[c]oming to understand has a particular phenomenological appearance” (that of insights) consisting in felt novelty, sometimes surprising enough to provoke the “Aha!” feeling. (I return to the issue of creativity in Section 8.)

In order to see what is novel in Lynch’s view, it may help to contrast it with Bourget’s view. Two differences stand out: what it is for understanding to occur, and how they conceive of grasp.

What Lynch means in saying that coming to understand is “partly definitive” of the understanding that ensues is that, in addition to the epistemic description that needs to be satisfied if a state of mind is to qualify as a state of understanding, that state of mind is also necessarily acquired in a given act of coming to understand. This, presumably, stands in contrast to other manifestations of dispositional understanding once we’ve already acquired that understanding. Manifesting an already acquired understanding surely often happens. So Lynch’s point must be that such manifestations, as opposed to coming to understand something for the first time, are not “partly definitive” of the state of understanding.

In setting the stage, in Section 2, for the question of how understanding and conscious experience relate, I granted Bourget's claim that dispositional understanding (and grasp) consist in the disposition to occurrently understand (and grasp). If Lynch is right, not all such occurrences are constitutive of dispositional understanding – only the initial occurrences in which we come to understand something for the first time count for identifying the state of understanding that ensues.

Turn now to how Lynch and Bourget conceive of grasp. It's analytic that any state ensues at some time or other, so it must be important for Lynch that the moment of coming to understand is a moment when a certain *act* takes place, "the act of grasping" (Lynch 2017, p. 201).⁷⁷ Acts cannot be dispositional. So, on Lynch's view, *contra* Bourget, there is no such thing as dispositional grasp. So as to keep terminology uniform, I will continue to use the word "grasp" as Bourget does, and will use the phrase "coming to grasp" for what Lynch calls "grasp," by analogy with "coming to understand."

As for how understanding and conscious experience relate, Lynch claims that all acts of coming to understand have the phenomenal character of insights, viz., felt novelty, sometimes attended by a strong enough surprise as to be properly described as a "Eureka!" or "Aha!" feeling. If *all* acts of grasp involve "some level of insight," as Lynch contends, it would seem that all acts of coming to grasp (and understand) have the phenomenal character of insights. If so, this would support to Bourget's claim that having a conscious experience with the content grasped is a necessary condition for grasping that

⁷⁷ This agential aspect of understanding, Lynch argues on p. 206, makes understanding "difficult or impossible to convey *via* testimony," because, however a good and discerning listener one may be, the epistemic value of testimony is often due to the credibility of its source and the evidence available in its support rather than anything pertaining specifically to the listener. Lynch seems to be thinking of understanding as a manifestation of epistemic *agency*, and of thinkers as consequently responsible for understanding. In this text, I bracket issues surrounding epistemic agency.

content.⁷⁸ Going beyond Bourget's claim, Lynch thinks that experiences of coming to understand share the phenomenal character of felt novelty (although how vivid the surprise we experience is will, of course, vary).

The question I pursue in discussing Lynch's view is how coming to understand relates to what it is like to undergo insights. In the next section, I give reasons for thinking we may come to understand something even if the attendant experience is not one of insight.

7. Intuition and insight

Contra Lynch's view that coming to understanding may only happen accompanied by "some level of insight," in this section I refer to a possible case where coming to understand *isn't* accompanied by insight, suggesting the burden of proof lies with Lynch to re-interpret the example I offer.

In Section 9, I'll present my own view of how conscious experience and understanding relate. I don't deny that understanding may – perhaps even always – be accompanied by some form of insight.

⁷⁸ Lynch also writes that "we can understand why something is the case without being consciously aware of that fact, and without the understanding being available for conscious attention. Your understanding in such a case is tacit or implicit. Arguably, much of what we understand we understand in this way, and we can forget that we understand something that we do understand" (Lynch 2017, p. 200). It seems reasonable to suppose that, when the stakes are low, we may also tacitly *come to* understand, not only manifest an understanding already acquired. Does this contradict Lynch's earlier claim that all coming to understand involve "some level of insight?" It doesn't, if we distinguish having "a particular phenomenological appearance" from "being available for conscious attention." As Lynch writes, "[s]omething can be available to conscious attention without, at that moment, being attended to" (footnote 9). It makes sense to suppose that, if what insights are like partly consists in being surprised at what one learns, some surprises may be faint enough for us to not pay any notice to them.

However, I don't presuppose it either. If the considerations of this section are on the right track, to claim that insight always accompanies coming to understand is highly controversial. By not making a controversial claim, my view is closer to plausibility, and we have a *pro tanto* reason for endorsing it over Lynch's view.

Let's start from an example:

For you to have an intuition that A is just for it to *seem* to you that A. Here 'seems' is understood, not as a cautionary or "hedging" term, but in its use as a term for a genuine kind of conscious episode. For example, when you first consider one of de Morgan's laws, often it neither seems to be true nor seems false; after a moment's reflection, however, something new happens: suddenly it just *seems* true. (Bealer 2000, p. 3)

Suppose Insa learns de Morgan's laws much as Bealer describes the episode. In coming to understand de Morgan's laws, the laws seem natural to Insa. They seem true to her, and she finds herself to be familiar with solving logic problems involving them, and with explaining the laws to others.⁷⁹ Only, in coming to understand the laws, Insa experiences no felt *novelty*. The laws seemed true to Insa from the very start, she didn't feel in the least bit surprised to learn them, and can now apply them seamlessly.

By Lynch's description of what insights are, and Bealer's description of what intuitions are, we

⁷⁹ This is only an example; we shouldn't conclude that intuitions always make the conceptions intuited clearer. We may, at times, for want of reflective scrutiny, find intuitive what is in fact muddled. Ewing writes that, for the everyman, "confused non-inferential cognitions form the basis, for instance, of the categories of substance and cause, not to mention the part they play in ethics and religion. Anyone who claims that the validity of the categories of cause and substance is self-evident is liable to be challenged to give a precise account of the propositions about them which he holds are immediately known, and he may find that any tenable propositions turn out to be so complicated that it is quite unreasonable to suppose them self-evident" (Ewing 1941, p. 25).

should say that, in coming to understand de Morgan's laws, Insa had an intuition that the laws are correct but no insight in learning them.⁸⁰ A case like Insa's is certainly conceivable. That gives us a *pro tanto* reason to think a case like Insa's is metaphysically possible. If it is possible, Insa's case is a counterexample to Lynch's claim that all instances of coming to understand are necessarily accompanied by "some level of insight."

In reply, Lynch may argue that Insa must have experienced an insight with its characteristic felt novelty, only she didn't notice it, and so her case is no counterexample. Call the case, so re-described, Elsa's case. Enough of us surely find ourselves in Insa or Elsa's position when first learning elementary logical laws: they hold no noticeable surprises, and finding them only puts order in what we previously thought. However, this reply only shows that cases like Elsa's are possible – *not* that cases like Insa's are impossible. Unless provided with an argument for why we cannot come to understand something without having an insight, Insa's case as initially described seems possible.

Incidentally, given that insights and intuitions are phenomenally distinct, if cases like Insa's are possible, we have a *pro tanto* reason to think that there need be no unique common phenomenal character to experiences of understanding regardless of what is understood. This points in the same direction as Bourget's view that what grasping a content is like varies with the object of the grasp, for different contents are presumably experienced differently.

In this section, I argued that coming to understand doesn't require conscious insights. I described

⁸⁰ I set up the case such that Insa should have an intuition in coming to understand de Morgan's laws – a conscious intellectual seeming – so that it should be clear that this counterexample to Lynch's view isn't a counterexample to Bourget's view as well. Note, incidentally, that not *all* counterexamples to Bourget's view are also counterexamples to Lynch's view. For perhaps there are cases of unconscious occurrent grasp in which we manifest an understanding we already possess, but acquiring that understanding couldn't occur without an initial insight.

Insa's case, in which coming to understand de Morgan's laws may be unaccompanied by either surprise or felt novelty. That case seems metaphysically possible, and it is a counterexample to Lynch's view.

8. Creativity and understanding

In this section, I present (and object to) Lynch's argument in favor of the claim that coming to understand requires insight. Lynch appeals to the thought that coming to understand is creative:

[C]oming to understand seems to fit this model of creative psychological acts: it involves generating in a synthetically imaginative way new and valuable ideas. Which ideas? Those that concern dependency relationships—how things fit together. The “grasping” of those relationships, which lies at the heart of understanding, is what makes understanding creative. (Lynch 2017, p. 202)

We might be tempted to deny that coming to understand is creative. What is at stake, however, is psychological or mental creativity, resulting in ideas that are new relative to old ones. And it's uncontroversial that coming to understand something changes some of our previous ideas about it, or leads us to form new ones.

For Lynch, “[h]aving... an insight is part of why understanding is fundamentally a creative act” (p. 202). Why so?

[I]t is worth emphasizing that the surprising... aspect of creativity makes creating seem at once something we do (which it is) and at the same time something happening to us.

The muse suddenly strikes. Realization comes in a flash. Coming to understand is like

this as well. It involves insight[.] (p. 204)

The thought seems to be that surprise correlates with the novelty of the ideas in whose production the thinker is creative, and that this applies, *a fortiori*, to the novelty of ideas had in coming to understand something new.

If so, Lynch seems to be proposing the following argument. If coming to understand is a creative act, then it requires the felt novelty and surprise characteristic of insight. And coming to understand *is* a creative act. It follows that coming to understand requires “some level” of insight.

The argument is valid. However, I believe that a dilemma can be raised against it. To do so, we need to ponder alternative views of what creativity might consist in. For instance, Nanay (2014) advocates a view on which

it is a necessary feature for creative mental processes that their outcome is experienced in a certain way: that we experience the outcome of the mental process as something we have not taken to be possible before. (Nanay 2014, p. 23)

Nanay doesn't think this *excludes* a functional view of creativity: “I am not denying that for each token creative process, there is (or at least can be) a functional/computational process that implements it” (p. 22). However, it is a familiar point that, when it comes to insights, we lack principles detailing how and when various functional processes with epistemic purport are accompanied by conscious experience. And this makes room for the following dilemma.

Suppose, on the one hand, that, *pro* Nanay, creativity is best accounted for in terms of conscious experiences. Both Nanay and Lynch consider surprise and “Aha!” experiences as paradigm experiences had in being creative. However, how are we, on purely phenomenological grounds, to argue that

understanding is creative? Recall Lynch's description of understanding quoted above:

its content concerns dependency relations between propositions or states of affairs relevant to the subject; it is conducive of the agent's ability to offer justified explanations of the relevant subject; and it disposes the agent to make further justified inferences both factual and counterfactual about the subject. (p. 199)

Nothing in this description guarantees that one's coming to understand something new needs to be accompanied by conscious surprise. Lynch says elsewhere:

Creative ideas are valuable to the person's cognitive workspace. They move things forward on the conceptual field on which they are currently playing. They are useful and fecund. They have progeny, and they contribute to the problems at hand. (p. 202)

This surely fits the most famous examples of scientific and artistic understanding. And, given that creative ideas "are valuable to the person's cognitive workspace," it would make sense that the person in question should be surprised in finding them. However, if we compare the descriptions of understanding and creative ideas in the two quotes above, it seems clearly conceivable that understanding may not *always* be creative in the more demanding sense suggested by the metaphors of fecundity and progeny.

Suppose, on the other hand, that, *contra* Nanay, creativity, at least in matters intellectual, is best accounted for in terms of the net epistemic result produced – not in terms of what experiences may accompany the process of producing those ideas. Lynch notes (p. 199) that understanding involves the agent's *ability* to justify the explanations they *can* give, and that understanding "disposes the agent to make further inferences." These are all dispositional terms that relate to conscious experience only indirectly if at all. On a purely epistemic description of what a creative act is, and of understanding, and granting that understanding is creative, while many experiences may accompany coming to understand,

it is unclear why any insight would be required.⁸¹

It seems that, whether we choose a purely experiential account of creativity, or a purely epistemic account of creativity, one of the two premises of Lynch's argument will be questionable. For it will turn out that either coming to understand isn't creative in the required sense, or that, if it is, that in itself doesn't require that those who understand should undergo any conscious insights.

In this section, I explored Lynch's argument for thinking that coming to understand does require conscious insight, and proposed a dilemma against it based on how we may conceive of creativity.

9. Modes of understanding

In this section, I present and argue in favor of my own conception of how understanding and conscious experience are related.

The following seem to be uncontroversial starting assumptions to make about understanding. Understanding needs to be *acquired* for it to be *possessed*, and it needs to be possessed in order to be *exploited*. An already possessed understanding is always exploited on a specific occasion. The occurrent

⁸¹ Lynch may reply by pointing out that, on his description, a creative act "involves generating in a synthetically imaginative way new and valuable ideas" (p. 202). And synthetic imagination, one may think, cannot but require conscious experience. However, Lynch has offered no account of imagination such that imagination should require conscious experience. For a simple example, imagine a unicorn. Surely what you imagined is a horse with a horn on its forehead. You probably didn't imagine the process of, say, gluing the horn to the horse's forehead. So, in the phrase "synthetic imagination," the synthesis is often not itself consciously experienced. Whereas what makes imagination creative, presumably, is precisely the production of "new and valuable ideas." So it would seem that, even in some of the cases where creativity and conscious experiences are correlated, the correlation isn't epistemically that significant.

state of mind manifesting understanding may have a discursive, propositional content. Or understanding may be exploited in fluent performance whose content may not be propositional in form.

Correspondingly, on my view, there are three modes in which grasp and understanding may occur. The first mode is coming to understand something for the first time. The second mode is one of discursively exploiting an understanding we already possess. And the third mode in which understanding may occur is non-discursively exploiting an understanding we possess. These three modes in which understanding may occur will matter in ascertaining the relationship between understanding and experience.

A clarification about the discursive/ non-discursive distinction. In characterizing grasp, Bourget writes:

We can grasp entities of many different kinds, for example, facts, propositions, concepts, definitions, theories, structures, processes, and phenomena. Here I focus on graspings of propositions. While I focus on graspings of propositions, it seems reasonable to hope that all other kinds of grasping might ultimately be reducible to graspings of propositions [...] (Bourget 2017, p. 287)

If the objects of understanding must only have a “proposition-like” structure (p. 304), the meaning of the word “proposition” seems quite watered down. It seems better to reserve the word “proposition” for what sentences express in their context of use (in language or thought), and to admit that grasp and understanding might have non-propositional objects in this narrower sense. (I prefer the phrase “non-discursive” to leave it open whether Bourget is right that all objects of understanding are “proposition-like.”)

If a dispositional understanding we already possess may be exploited in a non-discursive mode,

it is plausible to think that the existence of such a mode does justice to what has been variously called “ineffable” or “tacit understanding:” to instances of understanding that aren’t subject to an “articulacy requirement.”⁸² If one *denies* that instances of tacit understanding exist, one would shoulder the burden of proof to re-interpret instances of seemingly tacit understanding.

Perhaps Bourget is right to think that grasp may only occur consciously. But I have considered cases of unconscious grasp that are conceivable and, we thereby have a *pro tanto* reason to think, also metaphysically possible. And I have criticized the main argument that Bourget offers in favor of his view. So it seems more plausible to stay neutral on whether unconscious grasp is possible or not. On my view, it's not a requirement on grasp and understanding that they should occur consciously.⁸³

Conscious experiences that accompany understanding are quite varied. We sometimes understand at one blow. Otherwise what we understand dawns on us slowly, as the toil of inquiry comes to fruition. We sometimes hit upon an explanation by sheer chance. At other times, the sure hand and the quick eye leave no room for doubt we've understood what we're doing. At times, in understanding something, we see things differently than before. Often, understanding comes as the resolution of a puzzle we previously

⁸² The phrase “articulacy requirement” belongs to Grimm (2016). Lipton (2009) considers cases of tacit scientific understanding manifested, not by discursive knowledge, but by one’s cognitive abilities in manipulating the models one works with. Moore (1998) thinks “ineffable understanding” is a suitable way to characterize how the early Wittgenstein conceptualizes the cognitive contact we have with the arts, “whereof one cannot speak.”

⁸³ A similar remark applies to Lynch’s view. Lynch argues that coming to understand something for the first time requires conscious insight. I have considered cases of coming to understand that don’t seem to be accompanied by conscious insight. Such cases are obviously conceivable, and that gives us a *pro tanto* reason to think they are also metaphysically possible. And I have objected to the argument Lynch offered in favor of his view. So it would seem more likely to stay neutral on the question of whether coming to understand is possible in the absence of insight. On my view, insight (with its attendant phenomenal character of felt novelty and surprise) is not required in order to come to understand something new.

hadn't been able to solve. This variety in experiences of understanding is partly explained by Bourget's suggestion that what it is like to grasp different phenomena will be different according to what those phenomena are.⁸⁴

Variety of experiences notwithstanding, we have been given no decisive reason to think that a useful classification of experiences associated with understanding *cannot* be undertaken. One may plausibly think that many experiences are just fellow travelers with understanding without having any necessary connections with it. Suppose we grasp the meaning of a hieroglyph in a flash of insight. We're enthused by it, and feel deep relief that we finally solved the puzzle. It's surely possible that we may have understood the hieroglyph *without* feeling enthusiasm or relief, or while feeling them to a lesser extent than we did. Sometimes, sheer conscious grasp of content doesn't involve any feelings like enthusiasm or relief.

The question is whether such a line of thought could be extended to *all* experiences associated with grasp and understanding. We have, it seems to me, no decisive reason to think so. This is because the notion of a conscious grasp or apprehension of content that is accompanied by *no* conscious attitude with respect to it, and no feeling elicited by it, seems to be conceptually possible but inadequate to our cognitive make-up. *Can we* simply grasp something while taking no attitude whatsoever towards it – not even one of wavering hesitation? Once a question arises in conscious thought, if we cannot settle it one way or another, we seem to adopt *other* attitudes and feelings toward it: disbelief, doubt, confusion, or uncertainty. This surely is no conceptual necessity: but it seems closer to how we think than the possibility

⁸⁴ As Bourget argues, this is so because oftentimes (he thinks always), in grasp, we consciously experience the grasped content. And, if what an experience is like is partly determined by its content, it follows that what it is like to grasp may largely depend on what is grasped.

of bare, non-attitudinal content apprehensions unaccompanied by any feeling.⁸⁵

These considerations initially motivate – without making it compelling – the thought that we should be able to draw a distinction between conscious experiences *surrounding* understanding and conscious experiences *of* understanding, or experiences that understanding partly *consists in*.⁸⁶ Once we consider such a distinction, the question arises as to how to best draw it.

It's important not to draw the distinction between experiences of understanding and experiences surrounding understanding too narrowly. For instance, Lynch argues that coming to understand is conscious in moments of insight, whose phenomenal character partly consists in felt novelty and surprise. Moreover, he argues that coming to understand is “partly definitive” of the state of understanding that ensues. (Section 6 gives the details of his view.) It follows that conscious occurrences of an understanding already possessed don't define it, and don't seem to contribute anything constitutive: they merely *surround* the understanding manifested. Once the variety of experiences that are properly *of* understanding is limited to insights, it follows that the phenomenal character of insights – their felt novelty and surprise – will be typical of *all* experiences constitutive of understanding. I argued against this view in Sections 7-8, and it seems independently implausible. (I don't *deny* this implication of Lynch's view. However, since my own view stays neutral about the controversial implication that all

⁸⁵ Thagard (2002) argues in favor of the view that “hot cognition” – thinking accompanied by feeling – is prevalent in many scientific traditions, and doubts the possibility of emotionally disengaged thought.

⁸⁶ It isn't clear that such a distinction can ultimately be drawn, since to do so would involve showing how experiences of understanding play the epistemic roles attributed to understanding, such as that “its content concerns dependency relations between propositions or states of affairs relevant to the subject; it is conducive of the agent's ability to offer justified explanations of the relevant subject; and it disposes the agent to make further justified inferences both factual and counterfactual about the subject” (Lynch 2017, p. 199). I discuss this issue in Chapter IV.

experiences constitutive of understanding share what it is like to undergo them, my view is less controversial than Lynch's view.)

On the view I advocate, the kinds of conscious experience that seem to more closely align, in beings like us, to the three modes in which understanding may occur consciously are insights, intuitions, and experiences of fluency, respectively. *If* the occurrence of understanding partly consists in any experiences at all, my view provides that it partly consists in experiences that fall under one of these three categories whenever understanding occurs consciously.

- (a) Coming to understand is conscious in moments of insight. We experience insights when, in solving a problem, we suddenly see a correct solution to it, typically after casting the problem in different terms. These realizations are often called "Aha!" or "Eureka" moments.
- (b) Exploiting an understanding we already possess on a specific occasion, in a discursively articulate way, and without changing what we thought beforehand, gives rise to intellectual intuitions.
- (c) Finally, possessing a running understanding of something is conscious in experiences of fluency. We achieve fluent cognitive performance when we meet with no obstacle blocking cognitive success because we pay attention to the right things as we proceed.

There seem to be broad phenomenal characters associated with each category. As Lynch notes, what insights are like partly consists in the felt novelty of what one grasps and the surprise had in grasping it. Notice that whether coming to grasp something for the first time is propositional or not isn't decided by feelings of novelty or surprise. So there is no need to draw a distinction between discursive and tacit coming to understand.

Bourget, in turn, notes that the “perceptualist theory of intuition is the view that comes closest to the phenomenal theory of grasping” (p. 296). Whether intuitions should be best construed as a kind of perception or a kind of belief isn’t crucial here.⁸⁷ On both approaches, what is grasped seems to be true, and this seeming veridicality is part of what intuitions are like. Only propositions are apt to be true, and only grasped propositions are apt to seem to be true to the thinker who grasps them. So it’s plausible to think that intuitions are the kind of conscious experience that discursively (or propositionally) manifests an understanding that the thinker already possesses.

As for experiences of fluency, they often occur as part of what Nakamura and Csikszentmihalyi (2002) call “experiences of flow,” exemplified by feeling sure of what one does (or thinks) in agreement with one’s previous understanding of the situation. This felt continued assuredness that no mishap upsets one’s thinking or doing as it unfolds doesn’t presuppose that what we grasp is discursively explicit. Experiences of fluency are, then, the kind of experience that may manifest tacit understanding.

In saying that broad phenomenal characters are associated with each of insights, intuitions, and experiences of fluency, I’m not committed to saying that the phenomenal characters of these three kinds of experience are everything that conscious experiences of understanding consist in. What individual experiences are like may far outstrip what, in their phenomenal characters, may lead to classifying them as either insights, or intuitions, or experiences of fluency.

Insights, intuitions and experiences of fluency may co-occur.⁸⁸ It may be that, in solving a

⁸⁷ For defenders of each approach to intuitions, see Chudnoff (2013), and Bealer (2000).

⁸⁸ Co-occurrence does not efface the distinction between modes of understanding. One may argue that the first time one exploits an understanding one has is in its very formation. And exploiting that understanding may be done with ease and fluency.

problem that has been on your mind for a while, you have an experience of fluency, feeling you are making progress in exploring ways to conceptualize it. When finding the appropriate terms to cast the problem in, you may experience an insight, with its felt novelty and slight surprise. In that insight, you may find the solution you were working toward intuitive. If re-conceptualizing the problem and the particular solution you reach contribute to your better understanding of the problem, then insight, intuition and fluency will have co-occurred as faces of your overall experience of understanding.

I have advanced a new view concerning how understanding and conscious experience are related, distinguished three modes in which understanding may occur, and identified three kinds of conscious experience corresponding to the three modes of understanding. And I have argued that my view avoids making the questionable assumptions made by Bourget's and Lynch's competing views.

10. Conclusion

The view of how understanding and conscious experience are related that I advanced is the following. Understanding may occur in one of three modes, and a type of experience corresponds to each – inasmuch as understanding occurs consciously at all. The first mode in which understanding occurs is coming to understand something for the first time. Coming to understand is experienced in insights. The second and third modes in which understanding occurs concern the exploitation of an understanding we already possess. We may exploit such understanding on a specific occasion, giving it a discursive, propositional structure; and we do so consciously in intuitions. Or we may exploit our dispositional understanding on a given occasion as the background for our ongoing (possibly non-discursive) cognitive activity when we are immersed into our epistemic environment; if understanding occurs consciously in

this mode, it does so in experiences of fluency.

The support I have offered in favor of this view was to consider two other views of how understanding and conscious experience are related, those advanced by Bourget (2017) and Lynch (2017), to argue that they make two controversial assumptions, and to explain how my view stays neutral on whether those assumptions are true. Staying neutral on controversial assumptions is a step toward plausibility, which offers a *pro tanto* reason to comparatively prefer my view of how conscious experience and understanding are related over Bourget's view and Lynch's view.

Chapter IV

Understanding: conscious experience and epistemic norms

We only understand, say, the position on a chessboard, if that position is intelligible, or makes sense to us. And, often enough, there is something it is like to grasp the patterns on the board. The question I tackle is how the epistemic and phenomenal aspects of understanding are related to each other.

In order to address this question, we need preliminary accounts in the epistemology and phenomenology of understanding. However, lest conscious experiences met the epistemic norms of understanding merely by coincidence, we would seem to also need to answer the question of why epistemic and phenomenal aspects of understanding are related the way they are.

In this text, I argue that we seem to lack the evidence needed to adequately support the claim that there is a metaphysically necessary relation between what it is like to experience understanding, on the one hand, and meeting the epistemic norms of understanding, on the other hand.

1. Introduction

When a chess master looks at a chessboard, she grasps the patterns on the board, playing fluently, having the right intuitions about which plans are feasible, and sometimes experiencing insights

concerning combinations that may destabilize the opponent's position. The player understands the position on the board. If queried, she may, upon reflection and all else equal, later articulate her understanding of the position in explaining that position to others. She understands the position as a whole, in light of a coherent body of knowledge of an entire subject-matter – chess openings, famous games, and so on. She may typically be able to evaluate her current performance in light of past games.

The player's understanding of a chess position has both phenomenal and epistemic aspects. Phenomenally, the player consciously experiences a grasp of various patterns in their interaction.⁸⁹ Epistemically, her conception about how the game is going manifests her chess expertise. It's natural to think that not all the chess-player's experiences simply happen to accompany her understanding of the position on the board.

Conscious experience may be related in sundry ways to the occurrence of understanding. We may sigh in relief at finally having understood how to solve a problem in the distribution of humanitarian aid. Or we may be ecstatic about the nearing prospect of understanding a difficult proof in higher set theory. It would seem that relief and enthusiasm, in these cases, merely accompany the present or future occurrence of understanding – perhaps by happenstance.

Notwithstanding, several authors claim that there is a necessary metaphysical relation between

⁸⁹ We don't always experience something distinctive in understanding. Think of your everyday first-grader, busy with addition and subtraction. There need be no special "Aha!" or intuition or anything else in the way of their stream of consciousness when they get it, and come to apply those operations properly. However, the question of how epistemic and phenomenal aspects of understanding relate concerns those occasions in which these aspects *co-occur*. The question presupposes they both obtain and asks why; pointing out that phenomenal and epistemic aspects of understanding don't always co-occur is beside the point.

some conscious experiences and the occurrence of understanding.⁹⁰ In this text, I focus on whether any metaphysically necessary relation obtains between conscious experiences and the epistemic norms understanding meets.⁹¹

Lest any and all conscious experiences should meet the epistemic norms of understanding merely by coincidence if and when they do so, it would seem that we also need to answer what I will call, in Section 3, “Wittgenstein's question:” the question of *why* phenomenal and epistemic aspects of understanding are related the way they are supposed to be (in each specific view of understanding).

In this text, I will argue that we seem to lack the evidence needed to adequately support the claim that there is a metaphysically necessary relation between what it is like to experience understanding, on the one hand, and meeting the epistemic norms of understanding, on the other hand. That is, we may ultimately be unable to satisfactorily answer Wittgenstein's question. I support this conclusion by discussing two instances of theories of how phenomenal and epistemic aspects of understanding are related.

⁹⁰ For instance, Lynch (2017, pp. 199-203) argues – in a view explored and criticized above, in Chapter III – that the act of coming to understand is “partly definitive” of the resulting state of understanding, that the act of understanding is creative, and that “[h]aving... an insight is part of why understanding is fundamentally a creative act,” where insights have a “characteristic phenomenological appearance,” that of felt novelty and surprise. On Lynch's view, it would seem that feeling the surprise and novelty of the conception one forms is metaphysically necessary in order for that conception to constitute understanding something. Similar remarks are made by Grimm (2008) and Chudnoff (2013). The details of each of these views are quite elaborate, and go beyond the scope of this text.

⁹¹ I stay neutral about precisely what those norms are. Here is an example. Describing the epistemic norms met by a conception by which understand something, Lynch writes that: “its content concerns dependency relations between propositions or states of affairs relevant to the subject; it is conducive of the agent’s ability to offer justified explanations of the relevant subject; and it disposes the agent to make further justified inferences both factual and counterfactual about the subject” (Lynch 2017, p. 199). Since I discuss his view in Chapter III, I bracket the details of Lynch's specific approach.

The first is the view proposed by Bengson (2015), which I discuss in Sections 2-4. On Bengson's epistemology of understanding, the conception by which we understand something is more or less accurate, complete and coherent – and we understand the better, the better our conception meets these epistemic norms. On Bengson's phenomenology of understanding, we understand something in virtue of having a conscious intuition, or intellectual seeming, whose content is that what is understood seems to be true.

However, I argue that it's unclear, in Bengson's view, how the fact that an intuition meets the epistemic norms it does is supposed to lead to the fact that the conception acquired by having that intuition is supposed to constitute understanding by having (enhanced) accuracy, completeness, or coherence – the epistemic norms Bengson identifies for understanding. This shows that Bengson's view doesn't satisfactorily answer Wittgenstein's question of why phenomenal and epistemic aspects of understanding are related as they supposedly are.

The second instance of a theoretical view of how phenomenal and epistemic aspects of understanding relate isn't held by any one author, but is the view which seems to me to be most promising in addressing Wittgenstein's question. I reason hypothetically: if the view that seems most promising in tackling that question doesn't succeed in giving it a satisfactory answer, then no such answer is likely to be forthcoming.

On the view I find most promising for addressing Wittgenstein's question, and that I articulate in Sections 5-6, insights, intuitions, and experiences of fluency are conscious experiences that have been systematically linked to the occurrence of understanding. And the use of models, the ability to explain target phenomena, and manifesting cognitive skills are possible epistemic norms governing understanding.

How do these conscious experiences contribute to meeting these epistemic norms of understanding? In Section 7, I offer my hypothesis: that we understand by attending. Sections 8-9 make the case that different experiences of understanding meet distinct epistemic norms partly because, in each type of experience, we consciously attend to the phenomena understood in a manner specific to that kind of experience.

However, in Section 10, I argue that the evidence available only supports the claim that phenomenal and epistemic aspects of understanding are related with natural necessity at best, and that leaves open the question of whether that natural necessity is grounded in a metaphysical relation of causation or constitution. As elaborated by Bengson (cf. Section 3), what I call “Wittgenstein’s question” concerns precisely what evidence we may have for thinking that a metaphysically strong relation obtains between phenomenal and epistemic aspects of understanding. So conceived, Wittgenstein’s question can’t be given a fully satisfactory answer. The conclusion of the chapter explores the consequences of this result.

2. Bengson on understanding

In this section, I give a brief overview of Bengson's (2015) view of understanding. I discuss Bengson's view of understanding because I believe it is one of the more careful views of understanding recently advanced, and would serve well as a foil in the discussion to follow. On Bengson's view, understanding a phenomenon consists in having a conception of that phenomenon that is – to some contextually sufficient degree – correct, complete, and internally and externally coherent

(“coalescent”).⁹² A conception:

is how one conceives or thinks, or is somehow inclined to think, of it. As I shall understand them, conceptions have a tripartite structure, consisting of:

- a mental state or attitude, conceiving-of;
- an object, or what the conception is a conception of; and
- a content, or how the conception characterizes its object. (Bengson 2015, p. 651)

Bengson (pp. 654-655) argues that understanding something only consists in correct, complete, and internally and externally coherent conceptions about it. How well we understand is determined by the degrees to which the conception by which we understand is correct, complete, and coherent.⁹³

How does conscious experience relate to understanding on Bengson's view? He answers: “intuition is a source of understanding” (Bengson 2015, p. 640). As he clarifies the notion of a source,

a source of understanding is the epistemic basis of understanding. It is not merely a background condition for understanding—what enables, facilitates, or promotes the latter—but that in virtue of which one comes to enjoy a state of understanding, as and when one does. (*ibidem*)

⁹² Internal “coalescence” of a conception is its internal coherence, and its external coalescence is its coherence with conceptions of other phenomena relevant to the phenomenon therein conceived. Bengson prefers the word “coalescence” so as not to presuppose any specific account of coherence. While I prefer the more mundane word, his caveat carries over.

⁹³ Trade-offs are possible. High coherence may lower the bar for the accuracy of understanding; and the other way around. And, as one person's understanding may be more coherent than accurate, whereas another's may include large disparities but also be partly formed by a largely accurate conception, these features of understanding may also account for faultless disagreements between individuals who understand the same phenomenon differently.

For Bengson, the following three conditions have to be met if⁹⁴ we understand in virtue of having an intuition:

1⁰ Having an intuition causes one to understand.⁹⁵

2⁰ Intuition is (a part of) “the epistemic basis of understanding;” it (partly) justifies the conception by which we understand.⁹⁶

3⁰ The fact that one understands after an intuition is explained by one's having that intuition.

Given the epistemic norms governing understanding, it follows from Bengson's view that intuition has to lead to a conception which meets one or several of the norms of enhancing correctness, completeness, or coherence. Bengson writes:

By ‘intuition’, I mean a conscious non-sensory mental state or event in which it strikes one that things are a certain way—one has the experience of seeing, or seeming to see, that things are that way, when one reflects on the matter. For example, setting oneself to address the question whether identity is transitive, it may consciously strike one—one sees, or seems to see—that the answer is affirmative. (Bengson 2015, p. 641)

⁹⁴ In a companion piece (Bengson 2015b), Bengson admits that, in some cases, conscious experience may not be required for coming to understand. So conditions 1⁰-3⁰ have to be met only if it is in virtue of some conscious experience that we do come to understand.

⁹⁵ By causing one to have a conception that constitutes an understanding of something I mean acquiring, sustaining, or modifying that conception.

⁹⁶ Surely constraint 2⁰ is often too strong read without the parentheticals. What we understand isn't typically justified solely on the basis of undergoing specific intuitions. For it is hard to see how intuiting alone could endow understanding with its special epistemic positive status that sets it apart from plain knowledge and merely justified beliefs. So I read 2⁰ as: “Intuition is *a part of* ‘the epistemic basis of understanding,’ i.e. it is *a part of* what justifies the conception that constitute understanding.”

After considering a handful of everyday cases illustrating this notion of intuition, Bengson writes that conceptions meet the norms of understanding in virtue of having originated in intuitions:

I submit that in all of these cases there is an epistemic shift, involving a switch to understanding, and that the shift occurs in virtue of an... intuition, from which it follows that... intuition is a source of understanding. (Bengson 2015, p. 648)⁹⁷

The question I address next is how it might be that intuitions could lead to the epistemic shift that coming to understand consists in.

3. Wittgenstein's question

In this section, I articulate the question of why – or in virtue of what – the epistemic and phenomenal aspects of understanding are related, if and when they do so. Wittgenstein writes:

For even supposing I had found something that happened in all those cases of understanding, - why should *it* be the understanding? (Wittgenstein 1958, par. 153)

I construe Wittgenstein's question⁹⁸ as follows. Grant that understanding a phenomenon consists in having

⁹⁷ The text excised includes the qualifier “innocent intuition.” It was eliminated, twice, because Bengson identifies as innocent those intuitions that aren’t had in our explaining phenomena we understand, and he thinks it is true that explanatory intuitions give rise to understanding as well. It follows that intuitions – either innocent or explanatory – are sources of understanding, which is how the edited quote above reads.

⁹⁸ Wittgenstein's view is much richer, including more ways to criticize the project I pursue. Here, I am only concerned with what can be gleaned from the passage quoted. The question it raises may receive a variety of interpretations. E.g.: Why should understanding obey epistemic norms at all, and why these ones? Why should the conceptions by which we understand be sourced in conscious episodes at all, and why in intuitions and not in other conscious episodes (e.g.,

a conception of it that is (to some context-relative satisfactory extent) correct, complete, and coherent. Also grant that, in each instance of understanding, the relevant conceptions were formed in having conscious intuitions. Wittgenstein's question is in virtue of what conscious experience relates to the epistemic aspects of understanding.⁹⁹

It's useful to clarify Wittgenstein's question by appeal to Bengson's notion of a source of understanding. The intuitions that understanding is sourced in have to satisfy 1⁰-3⁰. It is an empirical issue whether, when and how intuitions cause understanding (1⁰); so let's grant it and consider only those cases where some intuitive conscious experience causes us to understand something. Similarly, some justificatory role has often been attributed to intuition with respect to the conceptions it helps bring about; so let us grant 2⁰ as well. Wittgenstein's question – as construed here – pertains to 3⁰. Ontically, we may ask why it is we understand in virtue of having an intuition. Epistemically, we may ask why a given intuitive episode explains the fact that we thereby understand.

In granting 2⁰, we grant that intuitions provide conceptions with *prima facie* justification.¹⁰⁰ On

introspecting and finding we would be willing to bet on the truth of how we understand things)? However, I only pursue the interpretation of Wittgenstein's question articulated in the main text.

⁹⁹ One might think that it is just a brute fact about us that conscious experiences accompany the understanding experienced, and that this is no mere accident but part of our cognitive make-up. Grimm (2016) might be sympathetic to this view, endorsing the claim that experiences of understanding are reliable conditional upon proper cognitive function. However, to say we're just *built* that way wouldn't be to answer Wittgenstein, it would be to vindicate him. This is because it would give a biological answer to a question posed to philosophers. And, if the proper answer is biological, it goes to show Wittgenstein was right that the problem wasn't philosophical to begin with, and that thinking otherwise (as presumably Bengson does) would be a mistake. What I'm exploring in the text is a way to answer Wittgenstein with a philosophical toolkit. Whether that project succeeds is assessed in Section 10 below.

¹⁰⁰ That is, that intuitions are part of an epistemic basis for understanding, providing an intuitive, seeming-based, or presentational *prima facie* justification for the conceptions they occasion. I stay neutral about the details of such a justification. In particular, I stay neutral about whether what is at issue in assessing the positive epistemic status of

Bengson's view of understanding, the conceptions by which we understand are (to varying degrees) correct, complete and coherent. It seems *possible* that our conceptions be *prima facie* justified by seeming true to us even though they *aren't* in fact, on balance, complete, correct, or coherent enough. So, granting 3⁰ doesn't follow from granting 2⁰.¹⁰¹

Consider an example. Suppose that the only way you could make sense of the 1905 Revolution was in terms in the ineptitude of Nikolai II as a czar, and the cruel directives he gave to disperse protesters by lethal force. Then your history instructor tells you a bit more of the population dynamics at the time, e.g., how poor peasants had quickly become, within a generation, even poorer factory workers, denied the right to unionize and to negotiate wages and other protections in the workplace (Carr 1961). You will have corrected your old conception of the causal structure of the 1905 Revolution. And you will have completed your conception up to a satisfactory degree. Your views are now more coherent, because you can match what your instructor says with what you've read in novels depicting urban poverty in late 19th and early 20th century Russia. Suppose that finding out each of these things was a little epiphany for you, an intuition. And suppose that intuition presented things in a new way to you, and that you're *prima facie* justified in believing that things are as they seem to you. Yet, even if those improvements to your understanding of the 1905 Revolution were caused by intuitions, and even if your intuitions justify them,

intuitions is *prima facie* justification (the semblance of justification, where appearances may be deceiving) or the existence of a *pro tanto* reason which, absent defeaters, would in fact justify what one believes following that intuition.

¹⁰¹ If intuitive justification is construed as a species of *internal* justification (Feldman and Connee 2001), then the conceptions justified may be *inaccurate*. Similarly, suppose we take *doxastic* coherence as a yardstick for justification, and the thinker is *unaware* of defeaters to the conception that seems true to them; it's still possible to consider those thinkers doxastically justified. They would be justified by the fact that the intuition that lead them to understand is introspected, and what is understood appears to be coherent with the thinker's web of belief (Lehrer and Paxson 1969). I don't assume any specific view of justification in what follows; these are only illustrations.

the conceptions may not be sourced in intuitions, because having those intuitions doesn't explain why the resulting conceptions meet the epistemic norms they do (coherence, completeness, accuracy).

Bengson thinks that being able to offer a satisfactory answer to Wittgenstein's question is a criterion of adequacy for theoretical views of understanding. He writes:

A philosophically satisfactory defense of the idea that intuition is a source of understanding in a given case or set of cases should do more than simply find or assign a role for intuition in coming to understand in those cases; what is wanted, in addition, is a story, or account, of how intuition manages to play the role it is being assigned. Were we in possession of such a story, then we would have an explanation of how intuition... could play the role of a source of understanding. (p. 650)¹⁰²

Taking Bengson's view as a guide, let us provisionally agree that being able to answer Wittgenstein's question – in particular by accounting for 3^0 – is an adequacy criterion that should be met by theoretical views of understanding.

In the next section, I argue that Bengson's own view fails to live up to this standard. Because of the richness and conceptual nuance of Bengson's view, the fact that it fails to offer resources with which to fully answer Wittgenstein's question raises another question, of whether other views may fare better; and, if not, what the root of the problem might be.

¹⁰² In view of a uniform terminology, I speak of meeting epistemic norms where Bengson speaks of playing epistemic roles.

4. Intuitions and the epistemic norms of understanding

For Bengson, understanding a phenomenon consists in having a conception of it that is (to some degree that is satisfactory in context) accurate, coherent both in itself and with the thinker's other relevant conceptions, and somewhat complete or lacking in gaps. In this section, I argue that Bengson's view of understanding doesn't have the resources to offer a plausible answer to Wittgenstein's question, because it doesn't afford a plausible story about how intuition contributes to meeting the epistemic norms of understanding.

It seems sensible to shift focus from justification to epistemic norms.¹⁰³ Alston (2005, pp. 26, 166) advocates reorienting epistemology away from disputes about justification and toward discussing epistemic "desiderata" (norms or requirements) that confer conceptions their positive epistemic status.¹⁰⁴ Staying neutral about Alston's more radical claims about the non-existence of justificatory properties, I join him in shifting the discussion to consider the epistemic norms applicable to, and appropriate for,

¹⁰³ One might think that Bengson's view *does* draw a clear epistemic path from intuition to understanding. Intuition justifies the conception by which we understand. Any conception by which we understand something is justified *eo ipso*, as that conception is (to some suitable degree) accurate, complete and coherent. Therefore, having an intuition leads us to form a conception that is (to some suitable degree) accurate, complete and coherent.

One might reply by saying that the argument just given equivocates over "justification," referring to a presentational justification suitable to intuitions in the first premise, and switching to a notion of justification that depicts accuracy, coherence and completeness in the second premise. Bengson doesn't make this argument, and I don't attribute it to him. My objection is that Bengson's view doesn't seem to provide us with the resources to explain how the epistemic norms intuition satisfies enable that intuition to lead to an understanding that meets its own (presumably different) norms.

¹⁰⁴ Here, I don't argue that Alston *is* right about this. I only point to the fact that understanding has routinely been identified as having a positive epistemic status regardless of whether the conceptions by which we understand are justified or not. For many (e.g., Zagzebski 2001, Elgin 2004, Alston 2005), understanding is an aim of cognition in its own right.

understanding.

Intuitions might satisfy epistemic norms themselves. Intuitions may, for instance, give their experiencers *pro tanto* reasons to believe the contents intuited *because* those contents were intuited.¹⁰⁵ Granting that they do so, it is unclear how that translates into an answer to Wittgenstein's question. For it is unclear how what epistemic norms intuitions satisfy *qua* experiences could lead to the fact that the resulting conceptions by which we understand satisfy *their* respective epistemic norms.

Typically, understanding is taken to differ from plain knowledge or merely justified belief.¹⁰⁶ However, let us grant for the moment that understanding is a species of knowledge,¹⁰⁷ so that the norm

¹⁰⁵ This may obtain either because specific intuitive contents are *pro tanto* reasons for the contents of beliefs formed on the basis of the intuitions having those contents (Huemer 2007); or because the very intuiting, as an experience, gives thinkers a *pro tanto* reason to believe whatever happens to be intuited (Peacocke 1998). For a discussion of whether intuitions provide us with evidence in favor of what is intuited, cf. Conee (2013). I remain neutral on whether intuition has any such positive epistemic status.

¹⁰⁶ By plain knowledge and merely justified belief I mean the kind of knowledge and belief that aren't epistemically characterized further in an interesting way. Contrast, for instance, rote knowledge, or knowledge gained from uncorroborated testimony from an unknown but credible enough source, with understanding what it is that such a source says; or contrast an understanding of what we remember to rote knowledge.

To differentiate understanding from plain knowledge, consider a different example. We all learned arithmetic in grade school. After the first few lessons, we could all see – intuit – that $3+5=8$. But did that amount to a genuine understanding? As grade students, we lacked the mathematical sophistication to ask foundational questions. And, in order to truly understand arithmetic rather than merely know how to add and multiply numbers, etc., we should be able to appreciate the purport of such foundational questions. It is mysterious how the phenomenal character of an intuition that $3+5=8$ might epistemically correct or complete our understanding of that arithmetical truth, or how intuition alone may make that understanding cohere more, internally or with our larger web of belief.

¹⁰⁷ It is often thought that “understanding is a species of knowledge” (Grimm 2006). Knowledge is warranted and true. Warrant and truth, however, are the kinds of thing our conscious experiences are blind to. (This is why evil demon scenarios get an initial grip on most of us, however we then respond to them.) Knowledge has external conditions of success, unlike conscious experience. This might seem to undermine the project of looking for metaphysically necessary phenomenal-epistemic relations in understanding. There is a way to avert this worry: by focusing on the internal epistemic

of accuracy is satisfied. It doesn't follow that the epistemic norms satisfied by intuitions translate into the satisfaction of the norms of coherence and completeness, even if the resulting state *is* one of knowledge.¹⁰⁸ In other words, when we understand by intuiting, it's not necessarily whether we reap knowledge that counts; it's the properties that knowledge possesses – coherence and completeness, for Bengson – that qualify it as understanding.

Consider completeness, conceived commonsensically as having a somewhat unified and unwanting grasp, with no glaring misconceptions, no absent relevant knowledge, and the like. Why would having an intuition lead to completing the conceptions by which we understand? It may help to revert to the example of the 1905 Revolution. As 6th grade pupils, we might say: the Czar ordered that the crowd be shot at, that fueled unrest, and the ensuing bloodshed is the tragic result of this escalation on both sides. In learning this sequence of events, we get some modest intuitive grasp of what happened. Yet, as Tetlock and Belkin (1996) explain, we would then have a purely event-based historical understanding. A fuller and deeper understanding would also consider systemic factors, such as the workforce's economic condition in 1905, or bureaucratic inability to manage demographic changes. If we were to rest satisfied with the intuitive who-shot-whom event-based discourse, we might be more prone *to miss* factors that would considerably complete our understanding of the revolution. And,

aspects of understanding. *For the purpose of* approaching the question of how phenomenal and epistemic aspects of understanding relate, we should consider the aspects of understanding that are internal to the mind.

¹⁰⁸ Notice that what is at issue isn't *a priori* knowledge, evidentially supported by intuitions, in Bealer's (2000) sense. For, on Bealer's view, intuitional content was conceptual, and it could *only* thereby deliver *a priori* knowledge. Whereas we may understand sundry phenomena, and only relatively few of them in a purely *a priori* way. I am neutral on the question of whether only conceptual content may be intuited, or precisely how that claim, if true, is to be further articulated. For discussion, see Brogaard (2013), who argues that intuitions themselves are non-conceptual experiences, but that the content intuited cannot but be conceptualizable.

conversely, a fuller picture of the causes of the 1905 Revolution may lose in intuitiveness on that count.¹⁰⁹

If it fails to enhance meeting specific norms for understanding like completeness, perhaps intuition may instead lead to balancing the norms of accuracy, completeness and coherence with one another. In even our most carefully thought through decisions, we rarely explicitly calculate which trade-offs between our values are optimal. What we find intuitive may help us strike these balances.

However, I think that can only happen in some cases. Imagine an especially careful thinker, who is guided by the second-order norm to base her conceptions in evidence to the optimal degree of support (assuming one exists), but who is also conscientious enough to know that she may wrongly estimate what degree of support suffices for that end. (And suppose both conscientiousness and sufficient evidential support for one's conceptions are norms for understanding.) After observing her previous intuitions in some specific domain and finding them reliable, our thinker may be conflicted. On the one hand, her current intuition gives her a good reason to have the conception she does and halt inquiry. On the other hand, she is conscientious and knows that having second thoughts and doubting her own intuitions might reveal previously overlooked evidence. May the thinker stop inquiry? We lack a general answer, however reliable the thinker's intuitions may be. This suggests that conflicts between epistemic norms (here, getting no more than optimal evidence vs. conscientiousness) cannot always be settled by intuitiveness.¹¹⁰ Intuitions can't always help us choose acceptable trade-offs between epistemic norms.

¹⁰⁹ The suggestion here is that intuitiveness may perhaps amount to misleading evidence in favor of believing what we find intuitive. That wouldn't make it any less evidence (cf. Cargile 1995).

¹¹⁰ One might think that such a conflict cannot exist, since conscientiousness would be vicious if it persevered despite having optimal evidential support for one's conceptions. That may be so, but it is unclear by what *first*-personal criterion we may discern virtuous from vicious conscientiousness in such cases. For discussion of this virtue, see Montmarquet (1987). For a similar but broader conflict between coherence and evidential standards, see Worsnip (2018).

Even when everything goes well, epistemically speaking, and having an intuition causes us to understand, justifies the ensuing understanding and qualifies it as knowledge, it still doesn't follow that that intuition is a *source* of understanding, *per* 3⁰. Failing that, Bengson's view lacks the resources to satisfactorily answer Wittgenstein's question. Bengson himself had argued that answering that question is a criterion of adequacy for theoretical views of understanding. Tentatively, I conclude that Bengson's view fails his own standard.

5. Conscious experiences associated with understanding

In Sections 2-4, I argued that Bengson's view of understanding, despite its conceptual richness and nuance, fails to afford a satisfactory answer to Wittgenstein's question of why phenomenal and epistemic aspects of understanding may relate in any metaphysically robust way (if they are so related).

In what follows, I argue that even on what may seem to be the most promising attempt to address Wittgenstein's question, we fall short of giving a satisfactory answer to that question. The view I will consider isn't held by any single author. Rather, it is the view that seems to me to come closest to answering Wittgenstein's question in a satisfactory manner. If this view ultimately fails to achieve that end, it would seem that *no* satisfactory answer to Wittgenstein's question is likely to be forthcoming.

To tackle the problem of how epistemic and phenomenal aspects of understanding relate, we need some conception of both the epistemology and the phenomenology of understanding. In the next two sections, I will offer preliminary conceptions of both the epistemology and phenomenology of understanding. My strategy is to have as rich conceptions of each of these as possible. For weaker views

may be incomplete in articulating epistemic-phenomenal connections in understanding, giving rise to spurious connections or missing extant ones.

While I find them plausible, I don't presuppose that the preliminary phenomenology and epistemology of understanding I give below are true. Rather, they are stepping stones for reaching a sharper formulation of our problem, namely, how one might answer Wittgenstein's question. Consequently, one may evaluate what follows in a conditional vein.

In this section, I offer a preliminary phenomenology of understanding. I'll first give a brief inventory of experiences linked¹¹¹ with understanding, and then I'll explore how these experiences may be classified. For Bengson (2015, pp. 640-641), intuitions are a source of understanding. Three different phenomenal aspects seem to be involved in intuitions, as he conceives them:

- (i) We experience a phenomenal contrast: between what we thought before intuiting (failing to understand something), and what we think now, after having the intuition (in understanding): "it strikes one that things are a certain way"
- (ii) We grasp (or manifestly understand) a content intuited: "one has the experience of seeing, or seeming to see, that things are that way"
- (iii) We experience seamlessly considering or entertaining in thought the phenomenon

¹¹¹ Our overall conscious life is quite rich in conscious episodes wherein we manifest an understanding of something. At times, we experience enthusiasm at solving a problem, disappointment at failing to do so, puzzlement in seeking a solution, surprise as to what solves a problem, or feel absorbed by inquiry. On other occasions, when we understand notwithstanding, we may fail to have any feelings of the sort. Conscious experiences that merely accompany the manifestation of understanding differ from conscious experiences that understanding originates in. However, if Wittgenstein's question can't be given a satisfactory answer, understanding might be sourced in *no* experiences whatsoever. In order to consider the issue, we first need a list of experiences systematically linked with understanding.

understood, with ease and familiarity.¹¹²

Call the experiences depicted in (i) “insights,” (ii) “intuitions” (narrowing Bengson's use), and (iii) “experiences of fluency.”

Insights, intuitions, and experiences of fluency are conscious experiences. There is something it is like to undergo them. And what it is like to undergo each differs from what it is like to undergo the others. So insights, intuitions, and experiences of fluency that manifest understanding¹¹³ are *separable* in conscious experience: we may have one without having the others.

That is, we may undergo an experience of fluency wherein we seem to find something intuitive although we cannot exactly put our finger on it. And we may have insights where we're not exactly sure what we've intuited yet we clearly experience the contrast insight brings (some revelatory experiences may be like that). And, of course, something may seem true to us without us undergoing any necessary contrast or fluency to go along.

Insights, intuitions, and experiences of fluency have often been identified as kinds of experiences that manifest understanding. Bengson makes this case for intuitions – conceived now along the lines of

¹¹² Bengson's (2015) work can nicely be factored into (i)-(iii). Bengson discusses “explanatory intuitions” as experiences of understanding. He also refers to the phenomenal contrast that accompanies the epistemic shift by which we come to understand. And he grants that the conceptions by which we understand something cannot always be put into words without remainder, being thus open to the possibility of tacit understanding. Intuitively experiencing such tacit understanding may, presumably, only be achieved in experiences of fluency, wherein one feels that things are going right, that one understands the situation as one should, yet one cannot fully articulate precisely what one's understanding consists in.

¹¹³ Many intuitions, insights and experiences of fluency don't manifest understanding; but I bracket those in what follows. False insights, a misleading sense of familiarity, or intuitions that rely on inaccurate conceptions fail to manifest genuine understanding.

(ii). Lynch makes the case for insights, discussed in Chapter III:

Coming to understand has a particular phenomenological appearance. In cases of sudden insight, this phenomenological aspect of creativity either constitutes, or leads to, the “eureka” feeling. But creative acts can be surprising even if they do not necessarily provoke that “aha” moment... Even when coming to understand happens gradually over time it still feels “new” – as if you couldn’t have understood it prior to that moment. It feels as if you’ve made forward progress. (Lynch 2017, p. 203)

Lynch argues that insights typify understanding: even if not all experiences of understanding are as phenomenally noticeable as “Aha!” experiences, all – he thinks – “involve some level of insight.” In insights (i), we experience a phenomenal contrast between one’s conceptions before having the insight – one in which one doesn’t understand a given target phenomenon – and one’s conceptions following the insight, having thereby come to understand that phenomenon.

As for (iii), understanding something also often involves experiences of *fluency*, in which our thinking proceeds seamlessly, correctly, and with the intended effect. Nakamura and Csikszentmihalyi describe experiences of flow wherein one is fluent at the cognitive task at hand, having

intense and focused concentration on what one is doing in the present moment, [a] loss of reflective self-consciousness [and] a sense that one can in principle deal with the situation because one knows how to respond to whatever happens next. (Nakamura and Csikszentmihalyi 2002, p. 90).

Doubtless, experiences of flow are complex. But part of them is the experienced fluency in solving the task at hand. How does fluency manifest understanding? Dieks and de Regt say:

A phenomenon P can be understood if a theory T of P exists that is intelligible (and meets the usual logical, methodological and empirical requirements).

A scientific theory T is intelligible for scientists (in context C) if they can recognise qualitatively characteristic consequences of T without performing exact calculations.

(Dieks and de Regt 2005, pp. 150, 151)

Why should recognizing qualitative consequences of a theory manifest understanding? As Dieks and de Regt write, “the present view rehabilitates an intuition... that understanding is related to familiarity” (p. 158), and recognizing the qualitative consequences of a theory indicates possessing an understanding of phenomena made intelligible by that theory. In other words, when an experience of fluency manifests one’s understanding of a phenomenon, fluency presupposes ease and familiarity in conceiving of the phenomenon in question.

Insights, intuitions, and experiences of fluency differ markedly in their phenomenal characters. Rather than single out one kind of experience as characteristic of understanding, I think we should endorse a healthy *pluralism* about the kinds of conscious experiences that manifest understanding.

Recognizing these different conscious experiences that have been systematically associated with understanding reconfigures Wittgenstein’s question, *since what it is like to understand may meet distinct epistemic norms according to what kind of experience manifests that understanding*. This will be important in ascertaining how phenomenal and epistemic aspects of understanding relate. In order to fully appreciate the point, we need to revisit the epistemology of understanding first.

6. Understanding, explanation, models and skills

At the start of Section 5, I laid out the strategy of offering an epistemology and a phenomenology of understanding that are as rich as possible, so as to explore the possibility of robust relationships between phenomenal and epistemic aspects of understanding. In this section, I apply that strategy by considering de Regt's rich epistemic description of understanding.

According to de Regt and Dieks (2005), we understand scientific phenomena only in light of theories we have of them, and only to the extent that we are familiar enough with those theories to recognize their qualitative consequences relevant to the problems we encounter. For de Regt and Dieks, we understand by devising and using specific cultural artifacts – theories, models, or metaphysical outlooks. As de Regt (2004) argues, researchers grow in expertise, becoming able to understand a wider range of phenomena to a greater extent, in interaction with what epistemic norms are met by the theories and models they use:

[W]hether scientists are able to use a theory for explaining a phenomenon depends both on their skills and on the virtues of the theory. More precisely, it depends on whether the right combination of scientists' skills and theoretical virtues is realized. Particular virtues of theories, e.g., visualizability or simplicity, may be valued by scientists because they facilitate the use of the theory in constructing models and predicting or explaining phenomena; in this sense they are pragmatic virtues. But not all scientists value the same qualities: their preferences are related to their skills, acquired by training and experience, and to other contextual factors such as their background knowledge, metaphysical commitments, and the properties of already entrenched theories.

...I define the *intelligibility* of a theory as the positive value that scientists attribute to the cluster of theoretical virtues that help them in their use of the theory.... scientists need intelligible theories in order to achieve scientific understanding of phenomena. (de Regt 2004, p. 103)

De Regt proposes this view as an account of scientific understanding, but I will consider it to hold more broadly for understanding in general. For our purposes here, three elements stand out.

- (iv) We understand a phenomenon only when we are able to explain it.
- (v) We understand a phenomenon only by bringing to bear models that represent that phenomenon (and what makes a difference to its occurrence).
- (vi) We understand a phenomenon only if we manifest our cognitive skills in conceiving it.

Each of these claims has received support in the literature, but also criticism;¹¹⁴ and they can be held jointly or separately.¹¹⁵ Internalizing models,¹¹⁶ affording explanations of what they conceive, and

¹¹⁴ For instance, explainability of the phenomena understood as a norm of understanding has been defended by Grimm (2006) and Strevens (2013). Model-use in understanding has been defended by Strevens (2013) and Wilkenfeld (2013). Manifesting expertise in understanding has been defended by Zagzebski (2001) and Hills (2015).

¹¹⁵ The richness of this epistemology of understanding suggests why it has so far been difficult to specify exactly what “objectual” understanding consists in, or what the “perspective” one understands from is. For views that rely on these notions, see Elgin (2004) and Giere (2010).

¹¹⁶ It isn't that I *endorse* de Regt's epistemology of understanding over Bengson's or other views. Rather, models are artifacts poised for public scrutiny. So they can be assessed for accuracy with respect to their target phenomena, for internal coherence, for depth of mathematical or intentional structure, and for the range of phenomena they are intended to capture – and thereby for completeness and coherence with our wider web of belief as well. Moreover, models are more readily assessed for how they satisfy these norms in a precise way, unlike conceptions. It is hard to see how one could *measure* how accurate or coherent conceptions are. Whereas models are not only open to evaluation; that evaluation can be made technically precise. (E.g., for probabilistic measures of the internal coherence of a model, see the discussion in Douven and Meijs 2007). In Bengson's terms, this suggests an epistemic norm for understanding: that the conception by which we

manifesting cognitive skills, are properties of the conceptions by which we understand.

De Regt's is a demanding¹¹⁷ list of epistemic norms for understanding. I take this to be an advantage. For a rich epistemic description of understanding should afford a substantial account of how epistemic and phenomenal aspects of understanding relate.¹¹⁸ However, each of (iv)-(vi) is plausibly subject to presumptive counterexamples, and interpretive work would be needed to argue that they in fact depict genuine epistemic norms for understanding. This is why, where de Regt thinks all (iv)-(vi) apply to understanding jointly, I will consider an epistemology of understanding on which the epistemic norm to be satisfied is specified by the disjunction of (iv)-(vi). We understand a phenomenon if either we are able to explain it, or we bring to bear models that represent that phenomenon (and what makes a difference to its occurrence), or when we manifest our cognitive skills in conceiving it.

understand something should be able to be translated into a publicly available model, and would, in that sense, be thought to partly internalize the model in question.

Our ability to in-principle measure coherence fits well a “possible public validation constraint” for understanding (Goldman 1997, p. 538). Such a constraint, in turn, seems plausible because it would do justice to a variety of norms floated in the literature: that we only understand a phenomenon if we can solve problems that involve it; that we only understand a phenomenon if we can produce (checkable) explanations of why it occurs; that we only understand a phenomenon if we can give reasons to potentially convince others to see it the way we do (showing our math); and that we are able to do all this, if we really understand, in a reliable way (see Khalifa 2012, Grimm 2016, and Newman 2017 for suggestions along these lines).

¹¹⁷ The list could include more items. Understanding may be responsive to reflective evaluation (Sosa 1997), or to self-knowledge (Zagzebski 2001). These further epistemic norms for understanding, however, are more controversial than the ones listed in the text. Importantly, I don't impose the constraint that the knowledge mobilized in understanding is know-that, but only know-how (in model-building and model-use). For discussion, see Wilkenfeld (2013).

¹¹⁸ If one rejects any of (iv)-(vi) as claims stating epistemic norms for understanding, then we simply obtain alternative notions of understanding, encompassing more conceptions as they require less of them.

7. We understand by attending

On the view articulated in the last two sections for the purpose of exploring how phenomenal and epistemic aspects of understanding are related, conscious experiences typically associated with understanding may be classified under (i)-(iii), and, epistemically, the conceptions by which we understand must meet one of the norms depicted by (iv)-(vi). The question, now, is how the experiences depicted by (i)-(iii) and the norms depicted by (iv)-(vi) relate.

Many possible relationships may obtain between insights, intuitions, and experiences of fluency, on the one hand, and the use of models of what we understand, being able to explain what we understand, and manifesting intellectual virtues or skills in understanding, on the other hand. However, to the extent that authors like Bengson (Sections 2-4 of this chapter) and Lynch (Chapter III) are right that a general metaphysically necessary relation should pair conscious experiences associated with understanding and the epistemic norms met by the understanding that is experienced, we should be a *common* element that in some sense unifies the disparate relationships we may identify across the board between (i)-(iii) and (iv)-(vi).

I will be arguing that, for the purpose of addressing Wittgenstein's question, the best candidate for what plays the mediating role between phenomenal and epistemic aspects of understanding is attention. In brief, *we understand by attending* to the phenomena understood.¹¹⁹ Attention has an

¹¹⁹ The idea that we understand by attending is indirectly supported by a view concerning what makes models have explanatory power. Ylikoski and Kuorikoski (2010) argue that a necessary condition for a model to possess explanatory power is that that model should be cognitively salient to the thinkers who use it to explain phenomena the model represents. In particular, this applies to the models we appeal to in understanding something, and to the explanations we may thereby provide for what is understood. The key notion here is that of cognitive salience, but that is as obscure as it is central.

epistemic side. I will argue that how we attend to what we understand contributes to how the conceptions by which we understand meet the epistemic norms of understanding. Attention also has a phenomenal side. I will argue that how we attend to what we understand varies in kind with each of the experiences of understanding considered: insights, intuitions, and experiences of fluency.

The thought that attention matters in understanding isn't new. The oldest clear statement I could find belongs to Pierre Laromiguière (1823, p. 312): “L' entendement est la réunion de l' attention, de la comparaison, et du raisonnement.”¹²⁰ However, Laromiguière seems to use “understanding” as a catch-all for most cognitive activity, in the tradition of the 17th and 18th centuries.¹²¹ Moreover, saying that attention is a part of understanding doesn't clarify what epistemic norms attention thereby meets. Although I see Laromiguière as a forerunner, how I develop the claim here is novel.¹²²

At issue here is *attention in thought*. As James (1890, p. 403) puts it: “Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem

Salience, however, may be construed as a property of what springs to mind because we attend to it. And this leads to the idea that, when our understanding is mediated by model-based explanations of what we understand, attention plays a key role in understanding.

¹²⁰ “The understanding comprises attention, comparison, and reasoning.”

¹²¹ Nor do I take a stance for or against views on which consciousness is a second-order mental state or process, especially given the 17th century view of conceiving consciousness that way, and of requiring that we may count as thinking at all – and, *a fortiori*, as understanding something – only if our thoughts were conscious. For discussion, see Jorgensen (2014).

¹²² I remain neutral about issues pertaining to the ultimate nature of attention. I also remain neutral about whether attention is a cognitive structure, process, or event. Nothing in the text hangs on making one choice over another. For instance, if paying attention structures the content attended to, or if attending is cognitive unison, nothing changes in my approach as stated in the text. That approach, however, does rule out views like Broadbent's, on which attention is essentially perceptual, for then it would seem to follow there is no attending in thought.

several simultaneously possible objects or trains of thought.” The solution I offer to the problem raised is conditional upon the existence of attention in thought. If understanding involves attention in thought, that is not to presuppose – nor deny – that attention is directed towards our thoughts themselves rather than what our thoughts are about; attention *in* thought need not be attention *to* thought.

It can't just be that *what* we attend to is what we understand. For worldly phenomena are the same whether they are the content of knowledge by testimony from a reliable source, a mere seeming, or of understanding full well. So these conceptions must also be differentiated with respect to *the manner* in which we attend to that content. I will argue that different kinds of experiences of understanding involve different manners of attending to what is understood. For each mode of understanding, I will propose an argument of the following form:

- (a) If we undergo an experience of type ϕ , then we attend in manner ψ to the content of that experience.
- (b) If we attend in manner ψ to a content, then we meet epistemic norm ρ .
- (c) So, if we undergo an experience of type ϕ , then we meet epistemic norm ρ .

The ϕ s are types of experiences of understanding depicted in (i)-(iii). The epistemic norms ρ are depicted in (iv)-(vi); they are norms that understanding meets. As the arguments below are of this form, (a)-(c) offers a *schema* by which to relate phenomenal and epistemic aspects of understanding; I present these arguments in Section 8, and mention sources of evidence that support the premises in Section 9.

However, depending on how we interpret the “if-then” conditional in (a)-(c), arguments of this form may or may not answer Wittgenstein's question by giving reasons for *why* experiences of understanding meet specific epistemic norms of understanding. I discuss the issue in Section 10, arguing

that arguments of this form are the best we can achieve, yet they fail to fully answer Wittgenstein's question by staying neutral between correlation, causation and constitution as possible metaphysical relations between phenomenal and epistemic aspects of understanding.

8. How are phenomenal and epistemic aspects of understanding related?

In this section, I show how (i)-(iii) may plausibly relate to (iv)-(vi): how consciously experiencing understanding may satisfy the epistemic norms of understanding. I discuss the premises of these arguments in the next section.

- (1a) If we have an insight wherein we come to understand a phenomenon, then we change our pattern of attending to that phenomenon so as to better discriminate what makes a difference to its occurrence.
- (1b) If we change our pattern of attending to a phenomenon so as to better discriminate what makes a difference to its occurrence, then we are enabled to explain that phenomenon.
- (1c) So, if we have an insight wherein we come to understand a phenomenon, then we are enabled to explain that phenomenon's occurrence.
- (2a) If we have an intuition that manifests an understanding we already possess of a phenomenon, then we focus our attention on that phenomenon.
- (2b) If we focus our attention on a phenomenon, then we use models to represent that phenomenon.

(2c) So, if we have an intuition that manifests an understanding we already possess of a phenomenon, then we use models to represent that phenomenon.

(3a) If we have an experience of fluency that manifests a tacit understanding of a phenomenon, then we attend to that phenomenon with ease and familiarity.

(3b) If we attend to a phenomenon with ease and familiarity, then we manifest cognitive skills in conceiving that phenomenon.

(3c) So, if we have an experience of fluency that manifests a tacit understanding of a phenomenon, then we manifest cognitive skills in conceiving that phenomenon.

The arguments are valid. In each claim, the consequent expresses a necessary condition for the antecedent. And the antecedent expresses a sufficient condition for the consequent. To wit: What it is like to undergo those specific experiences of understanding *wouldn't be what it is* unless these experiences met the epistemic norms that they do. And undergoing them is sufficient to indicate that we understand.

The conclusion of each argument *relates* a type of experience of understanding with an epistemic norm for understanding. We now have a *pro tanto* reason to think that undergoing experiences of understanding may partly explain why we understand in having such experiences.

Moreover, in each argument, the middle term is attention. That is, how we attend to what we understand mediates how insights, intuitions and experiences of fluency meet the epistemic norms of understanding. This supports the hypothesis I proposed: that we understand by attending.

What is the force of the presumptive necessitation codified by “if-then”? In order to satisfy condition 3⁰ of Bengson’s notion of a “source of understanding,” and thereby offer a fully metaphysically

satisfactory answer to Wittgenstein's question, the "if-then" in each of these claims should be that of metaphysical necessitation, affording, by the epistemic access we may have to it, an explanation of why it is we understand *in virtue of* experiencing insights, intuitions, or fluency.

Yet this is not what we find. In Section 9, I mention a few sources of evidence for the premises of the arguments just formulated. In Section 10, I argue that *the most* these sources of evidence support is the thought that insights, intuitions, and experiences of fluency meet the epistemic norms of understanding with *natural* necessity: as a matter of how we're built cognitively, as it were, and not as a matter of *metaphysical* necessity.

9. Defending the premises

In this section, I mention a few sources of evidence that make the premises of the arguments in Section 8 plausible.

Consider a research program in developmental psychology carried out by Gopnik and collaborators. She writes:

[T]here is a distinctive phenomenology associated with explanation. The phenomenology involves both the search for explanation and a recognition of the fact that an explanation has been reached. We might call them the 'hmm' and the 'aha'. In English they seem to be expressed by 'why?' and 'because'... This explanatory phenomenology also appears to be strikingly domain general. We seek and are satisfied by explanations of physical objects, animate beings, psychological agents, and even social groups. We seek and are

satisfied by explanations in terms of physical laws, biological processes, reasons, or rules. At least first-person, the aha of understanding why the wiring doesn't work seems quite similar to the aha of understanding why the electrician won't tell you why the wiring doesn't work. Even in children 'Why?' and 'Because' seem to cut across domains in this way. (Gopnik 1998, pp. 108-109)

In support of (1c), Gopnik thinks there is a constant phenomenal component to attaining the kind of understanding we have of what we can explain. When one finds the explanation one was looking for, the typical reaction is "Aha!" When one discovers there's something to be explained for which one currently lacks an explanation, the typical reaction is "Hmm.." In support of (1b), the shift from "Hmm.." to "Aha!" is the insight, had when we home in on what makes a difference to the target phenomenon, enabling us to explain it.¹²³

(1a) is supported by studies of insightful problem-solving. Ansburg and Hill (2003) identify the problems that require an insight in order to be solved as those problems where we first get stuck by focusing our conscious attention to what we shouldn't be if we were to solve the problem. But, they claim, we unconsciously attend¹²⁴ to what seem to us to be peripheral features of the problem-situation. By doing so, we notice – and this is the moment of conscious insight – some other features of the problem that enable us to provide a solution. On their view, we shift our attention so that we come to understand

¹²³ The idea of understanding on the basis of identifying what makes a difference to the phenomenon understood, so that we may be able to explain it, is also articulated by Strevens (2013).

¹²⁴ My claim that experiences of understanding are phenomenally unified by attention – when attending is conscious – doesn't imply that attending cannot always be conscious. Perhaps instances of subliminal attentional processing or unconscious attentional processing could be re-described. Only needed here is granting that the notion of unconscious attending is not incoherent.

what matters, and that shift is the insight.¹²⁵

(2a) says we focus on what we intuit. This, I submit, is trivially true. What might it possibly mean to say that we consciously intuit something, and that it strikes us as true, unless it partly meant that we focus our attention on it? More roundabout reasons in its favor exist,¹²⁶ but premise (2a) seems plausible enough on its own.

The connection between attentional focus and model-use in (2b) is supported by psychological studies of the *Einstellung* (or mental set) effect. As Bilalić et al. write,

Problem-solving failures caused by the *Einstellung* effect are the downside of a normally efficient cognitive mechanism. To deal quickly with the familiar, we rely on the knowledge acquired through past experience. It seems inefficient to spend time looking for an alternative solution if we already have an adequate one. Indeed, in complex real-world situations people usually prefer to look for solutions that are good enough rather than trying for an elusive best that may be out of reach [...] (Bilalić et al. 2010, p. 114)

When we are mentally set on a given solution to a problem (in chess, mathematics, etc.) it is difficult to change what we attend to precisely because what we focus on is determined by what models we use (including our web of domain-specific prior beliefs). As (2b) claims, the models we use to conceive of a

¹²⁵ As Ansburg and Hill admit, what counts as a problem and what counts as creativity are moot questions. This may apply, however, to some of Ansburg and Hill's own conclusions. Their dichotomy between analytical thinking and creative thinking in problem-solving, for instance, may be questioned by considering the important leap in understanding, affording *creativity* unseen heretofore, in Boole, Frege, and Russell and Whitehead's *analytical* formulations of propositional logic.

¹²⁶ Here is one such argument. We are acquainted with what we intuit. Acquaintance with a content involves mentally pointing at that content. But direct mental reference presupposes focusing our attention to what is referred to. So, intuiting a content presupposes focusing our attention to what is referred to.

phenomenon shape how we focus on it.

How do experiences of fluency, skill, and attention relate? Nakamura and Csikszentmihalyi (2002, p. 90) suggests that a key element of flow is “intense and focused concentration.” The ease and familiarity that are part of experiences of fluency must be due to this focus and concentration – properties of the manner in which we attend to what we are fluent in, and that lends support to (3a)-(3c). More recently, virtue-theoretical approaches to practical rationality have also offered considerations in support of (3a)-(3c).¹²⁷

Naturally, one may deny any of these premises.¹²⁸ Doing so would, however, still enable one to

¹²⁷ Railton writes: “Musical improvisation, whole-hearted engagement in a conversation or sport, and exercise of a skilled craft to solve a challenging problem are, I would argue, our paradigms of self-directed activity. We see the individual as agent especially clearly in such cases, for it is here that she has the greatest degree of control, self-expression, and attunement to available reasons. Moreover, we are told, individuals characteristically report such activities as optimal experience. Aristotle would explain: These activities show man acting in accord with his nature, making fullest use of his distinctive capacities. Having such experiences requires, I have argued, not just practical reason and will power, but practical competence, even practical fluency. What might a list of practical competencies look like? Here’s a start. Many general-purpose competencies figure among the conditions for practical competence – e.g. perceptual abilities, inferential capacities, multi-tasking capability, reliable memory, vivid imagination, etc. Others are especially relevant to action and emotional response” (Railton 2009, p. 108).

If by “fluency” we mean a complex conscious experience, it clearly couldn’t be what it is without having ease and familiarity as phenomenal components. (3a) seems close enough to a necessary truth not to require further support. Ease and familiarity, Railton suggests *pro* (3b), fallibly indicate that who experiences them is skilled enough at solving the task at hand (whether it be overt behavior, such as driving a car, or a cognitive activity).

¹²⁸ For instance, (2b) “If we focus our attention on a phenomenon, then we use models to represent that phenomenon” seems plausible *only* if one takes any of two theoretically unsavory options. The first option would be to expand the use of the word “model” so that it covers much of our mental representations that are coherent or substantive enough to be called “models” by any stretch of the word. The second option would be to restrict the use of the word “attention” so that, properly speaking (by some standard of propriety), cognitive mechanisms that might otherwise pass for attentional mechanisms fail that test because they involve no models (mental or external). Neither route seems fruitful.

follow the argument developed in Section 8 in a conditional vein.

10. Addressing Wittgenstein's question

If the epistemology and phenomenology sketched in Sections 5-6 are assumed, then one suitable way to relate phenomenal and epistemic aspects of understanding seems to be given by the following argument:¹²⁹

- (0) If we undergo a conscious experience that manifests an understanding of a phenomenon, then that experience is either an insight, or an intuition, or an experience of fluency.
- (1c) If we have an insight wherein we come to understand a phenomenon, then we are enabled to explain that phenomenon's occurrence.
- (2c) If we have an intuition that manifests an understanding we already possess of a phenomenon, then we use models to represent that phenomenon.
- (3c) If we have an experience of fluency that manifests a tacit understanding of a phenomenon, then we manifest cognitive skills in conceiving that phenomenon.
- (4) So, if we undergo a conscious experience that manifests an understanding of a phenomenon, then we are either enabled to explain that phenomenon, or we use models to represent that phenomenon, or we manifest cognitive skills in conceiving that phenomenon.

¹²⁹ The argument relies on sub-arguments in Section 8 and the sources of evidence in favor of the premises of those arguments mentioned in Section 9.

Conscious experiences associated with understanding lead to conceptions that meet different epistemic norms according to what kind of experience originated them. The argument supporting this conclusion offers a principled reason to think so, provided its premises are accepted.

I proposed that we understand by attending, and motivated this claim by considering both phenomenal and epistemic aspects of attending. It's hard to see how else we may answer Wittgenstein's question if we didn't pursue the strategy of identifying elements of understanding that have both phenomenal and epistemic aspects. And it's hard to see what better candidate than attention may relate the phenomenal and epistemic aspects of understanding so straightforwardly. So, while I don't endorse (4), I think it's the best approach to Wittgenstein's question: it shows how, in principle, we may relate phenomenal and epistemic aspects of understanding.

In a way, by the argument developed here, part of the sting is taken out of Wittgenstein's question. For the challenge of how the epistemic and phenomenal aspects of understanding relate is no longer an intractable *conceptual* difficulty. Instead, it is amenable to argument and to an assessment of the starting premises for plausibility.

We will, however, have failed to fully answer Wittgenstein's question, because the challenge was *not* to produce the science that may warrant claims of natural necessity to the effect that we understand *following* some experience. Rather, the challenge was to produce – or to acknowledge such a thing isn't possible – a metaphysical explanation of how it is that we manifest understanding *in virtue of*, or metaphysically necessitated by, undergoing some experience.

Now let's think of what the "if-then" dependencies represent in the premises and conclusions of the arguments developed in the last three sections, to see if claims of metaphysical necessitation are supported or not.

As Kozuch and Kriegel (2015) point out, when we consider how conscious experience may relate to *any* mental state, event or process, there are three possible kinds of metaphysical relations that may obtain: *correlation, causation, and constitution*. Even assuming that intuitions, insights and experiences of fluency systematically relate to understanding as a state with positive epistemic status, the question of *how* they relate is ultimately left open, as a matter of principle. Here are the three options we should consider:

(α) Experiences of understanding may systematically correlate with the understanding experienced without there being a further story – at a cognitive level, and with epistemic effects – to be told about how they do so.

(β) Or it might be that experiences of understanding are causally related to the understanding experienced in them. Insights, for instance, might causally *prompt* us to understand. And intuitions and experiences of fluency might be causal effects of an understanding previously possessed.

(γ) Or, respectively, it may be that insights, intuitions, and experiences of fluency *partly constitute* what meets the epistemic norms of understanding. On some occasions (when understanding occurs consciously), understanding *couldn't*, in a strong metaphysical sense, occur as it does, and with the epistemic properties it has, unless it occurred consciously.

The appeal to attention, and how what we attend to is presented to us in insights, intuitions, and experiences of fluency doesn't decide the issue between (α)-(γ). For the commonsense representations and empirical evidence we may bring to bear on the issue have *at best* the strength of *natural* necessity, staying neutral about whether that natural necessity is *metaphysically* grounded in causation, constitution, or sheer correlation.

11. Conclusion

How do phenomenal and epistemic aspects of understanding relate? To address this question, we need a phenomenology and an epistemology of understanding. In phenomenology, I identified insights, intuitions, and experiences of fluency as experiences of understanding. In epistemology, I identified the use of models, the ability to explain target phenomena, and manifesting cognitive skills as epistemic norms governing understanding.

Do these experiences relate to these norms for understanding? The answer I provided, (4), is that undergoing any one of these experiences of understanding suffices for meeting one or several of the norms for understanding. In other words, meeting one or another of these epistemic norms is necessary in order for insights, intuitions, or experiences of fluency, when we undergo them, to count as experiences of understanding.

And I have supported (4) by appealing to the phenomenal and epistemic aspects of attention, explaining how *it might be*, in principle (with a possible natural necessity, provided the premises of the argument I advanced are true), that the phenomenal aspects of understanding include attending to what is understood, and that the epistemic aspects of attending are part of why insights, intuitions, and experiences of fluency may contribute to the ensuing conceptions' meeting the norms of understanding.

While (4) addresses the question of how phenomenal and epistemic aspects of understanding relate, it fails to address Wittgenstein's question concerning the grounds of these phenomenal-epistemic relations. I delineated three equally well-supported yet mutually exclusive answers to this question: Either there is no interesting reason why phenomenal and epistemic aspects of understanding co-occur. Or being the type of experience that systematically correlates with understanding suffices to cause us to

understand. Or being a token experience of understanding suffices to partly constitute the understanding it manifests on that occasion. Between these options, I argued, we have no decisive reason to choose one over another.

Since what we were looking for was a metaphysically compelling answer to Wittgenstein's question, we seem to lack the evidence needed to adequately support the claim that there is a metaphysically necessary relation between what it is like to experience understanding, on the one hand, and meeting the epistemic norms of understanding, on the other hand.

I end by pointing to an important consequence of the failure to ultimately answer Wittgenstein's question satisfactorily. To the extent that a satisfactory answer to that question were deemed a criterion of adequacy for philosophical theories of how phenomenal and epistemic aspects of understanding are related (as we saw Bengson argue in Section 3), it would follow that there is a principled reason why contemporary theories of understanding have rarely accounted for how conscious experiences are related to meeting the epistemic norms of understanding. The reason is that claiming that any necessary relation exists between phenomenal and epistemic aspects of understanding cannot be adequately supported with evidence drawn from our own introspective experiences or cognitive-scientific projects exploring insights, intuitions, or experiences of fluency.

General conclusion

In conclusion, I will recapitulate the main points of each chapter, examine how they fit together, and say why the resulting outlook is important and what avenues are open for further research.

1. Summary

The question my dissertation tackles is how phenomenal and epistemic aspects of understanding relate. More specifically, the question concerns whether there is any metaphysically strong relationship between phenomenal and epistemic aspects of understanding. Supposing that some of our conscious experiences accompany understanding, does our understanding also partly consist in having these experiences, or is it partly caused by our having such experiences? Chapter IV answers this question, and the chapters that precede it offer preliminary accounts of the epistemology and phenomenology of understanding, respectively.

In Chapter I, I defended a virtue epistemology of understanding. The conceptions by which we understand something gain their positive epistemic status by being formed and accessed in the right way – that is, by manifesting cognitive skills. In Chapter II, I objected to Grimm's view of understanding, according to which we only understand a phenomenon if we can explain it, we may only explain it by knowing what it depends on, and we identify such dependencies by answering “What if things had been different?” questions. Jointly, these two chapters support a limited pluralism about what epistemic norms

states of understanding meet.

As for the phenomenology of understanding, Chapter III clarified the sense in which understanding may be consciously experienced, and defended the claim that understanding may occur in one of three modes. Either we come to understand something for the first time, or we exploit an understanding we already possess. If exploiting an understanding already at hand, either we do so in a propositional, discursive, conceptually structured way (this is the second mode of understanding), or we do so in a tacit, non-propositional manner (the third mode of understanding). And, I argued, three kinds of conscious experiences correspond to these three modes of understanding. Coming to understand is conscious in insights. Discursive exploitation of an understanding already possessed is conscious in intuitions, or intellectual seemings. And tacit understanding is manifest in experiences of fluency.

Following these preliminaries in the epistemology and phenomenology of understanding, in Chapter IV I argue that the evidence we possess for how phenomenal and epistemic aspects of understanding might, in principle, be related to each other at best supports claims we may advance with the modal strength of natural necessity. And that is consistent with our inability to adequately evidentially support claims made with the modal strength of metaphysical necessity. We are not, then, justified to say that meeting the epistemic norms of understanding may either partly consist in, or partly be causally explained by the conscious experiences that accompany the occurrence of understanding.

In this Conclusion, I explore a few consequences of the claims supported in previous chapters for several philosophical debates. I end by formulating a few open questions the dissertation raises, offering as many avenues for further research.

2. The unity of understanding

In the Introduction to this dissertation, I distinguished uses of the word “understanding” and its cognates, and homed in on what I have called “intellectual understanding.” One may still ask whether there is any genuine unity to understanding as depicted by *that* use. The remarks in the rest of this section partly address that challenge.

Consider the distinction between practical and theoretical understanding, which applies the traditional distinction between know-that and know-how. Without attempting to reduce either to the other, it follows from my dissertation that understanding might involve *both* know-how and know-that, and need not involve either. Here is why. On the one hand, the approach to understanding my dissertation takes is consistent with Elgin's (2004) view, on which understanding is not a species of knowledge at all, and false conceptions may constitute whatever propositionally-structured component our understanding may have (even though Elgin thinks know-how plays an important role in understanding). On the other hand, my approach is also consistent with Khalifa's (2016) approach, on which understanding a phenomenon consists in possessing (propositional) knowledge of what explains the occurrence of that phenomenon, not including elements of know-how as constitutive of understanding. While consistent with both Elgin and Khalifa's views, my view is also consistent with taking opposite views concerning know-that and know-how. Indeed, that follows from the pluralism about epistemic norms supported by Chapters I and II.

Or consider the distinction between objectual and propositional understanding (e.g., as drawn by Zagzebski 2001). It is often thought that one's grasp of an entire domain at a glance is properly called “objectual.” Whereas propositional understanding has a narrower object, discursively articulated as one

or several propositions. My approach to understanding allows us to make the best of both worlds. Models used in explanation need not be discursively articulated; they may be tacit. The explanations produced, however, may be discursively articulate. According to whether the use of models or the ability to explain what we understand are, respectively, taken to be required for understanding, that understanding may be either objectual or propositional. I remain neutral about whether these options exclude each other (cf. the discussion in Section 9 of Chapter III). For instance, if we only understand a phenomenon by being able to explain it based on models that represent what the occurrence of that phenomenon depends on, it seems natural to think that understanding is, to that extent, *both* objectual and propositional in different respects.

Consider, again, the distinction between subjective and objective understanding (e.g., as drawn by Trout 2002). Feelings of understanding are subjective. Chapter III addresses their existence, identifying insights, intuitions, and experiences of fluency as kinds of conscious experience that have been systematically linked to understanding. Objective understanding, it is often said, differs from such experiences. We gain objective understanding of a phenomenon when we are able to give it a purely ontic explanation. Feelings of understanding, it is contended, have nothing to do with this. Whether they occur is irrelevant, and may sometimes even hinder inquiry.

On the approach to understanding developed in Chapters I-II, we have no decisive reason to doubt that understanding is objective. To understand a phenomenon may involve being able to explain its occurrence, and the explanation in question might be ontic, rather than dependent on one or another model representing the phenomenon in question. However, Chapter IV suggests that the distinction between subjective and objective understanding may need to be re-thought. This is because Chapter IV is consistent with thinking that typical experiences in which understanding occurs may meet specific epistemic norms, and meeting such norms is a mark of objectivity even though having such experiences

is a mark of subjectivity.

In this section, I considered three distinctions: the distinction between know-how and know-that, that between objectual and propositional understanding, and the distinction between subjective and objective understanding. I have argued that drawing these distinctions does not undermine the project pursued in this dissertation, of exploring how phenomenal and epistemic aspects of understanding are related. On the contrary, it seems that results reached in the preceding chapters may lead us to re-think some of these distinctions; one future project is to explore how to re-think them.

3. Consequences of the dissertation for several philosophical debates

In this section, I consider a few consequences that the claims supported in this dissertation have for three philosophical debates pertaining to the nature of understanding or issues closely related with it. I briefly consider Meno's problem (in one acception), the relation between understanding and explanation, and the possibility of cognitive phenomenology.

Meno's problem. Recent epistemological interest in the nature of understanding was revived by considering states of mind, like understanding, that seem to be of special epistemic value to us as thinkers. In one version, Meno's problem – as dealt with by epistemologists (such as Pritchard 2008) rather than ancient philosophers – concerns the value of understanding: why understanding something is worth it.

The approach to understanding in this dissertation suggests – without implying it – that there may be several sources for the value of understanding. One is the value of knowing how to explain the phenomena understood. Another source of value is possessing and exploiting models of phenomena

understood. An additional source of value is given by manifesting cognitive skills in solving problems that involve phenomena we understand. Finally, there is the net benefit not just of truth, but of advancing our inquiries. We may often tap into several distinct sources for the value of understanding at once. This agrees with the pluralism concerning what epistemic norms properly apply to understanding that was supported by Chapters I and II.

Thus, the solution that this dissertation suggests for Meno's problem (in one form) is eclectic: there are several sources to the value of understanding. This affords a diagnosis for why Meno's problem has traditionally been deemed to be intractable, or at least conceptually intricate. That is because what was sought was a unique special kind of epistemic value (perhaps something in the vicinity of wisdom), whereas what was available in fact for showing what makes understanding valuable was only a mixed bag of subjective and objective epistemic values.

Conceptions of understanding and explanation. A second debate about the nature of understanding that the claims supported in this dissertation bear upon concerns how understanding a phenomenon and being able to explain that phenomenon are related. Is purchasing understanding the goal of explanatory inquiry, or is the verbal product of explanation the goal of our attempts to come to understand? Three main views have been proposed with respect to how to best conceptualize that relationship. Lipton (2009) has argued that there can be – and often is – “understanding without explanation.” In contrast, Strevens (2013) has argued that all understanding-why presupposes our ability to explain the phenomena understood. For both Lipton and Strevens, how understanding and explanation are related is a substantive philosophical issue; for Khalifa (2012), it is a verbal issue: we simply use the word “understanding” in relation to our explanatory abilities, and any novel inquiries into the metaphysics of understanding lack promise.

The pluralism concerning what epistemic norms states of understanding meet that was supported by Chapters I and II leaves it open whether being able to explain a phenomenon is a requirement for understanding it. In favor of Khalifa's view, one may choose to stipulate a set of norms that understanding satisfies, and to include explanation in that set. In the Introduction, I offered reasons both for and against choosing to stipulate the norms of understanding. However, *contra* Khalifa, what epistemic norms our states of mind meet, especially when we have good reasons for thinking those states qualify as states of understanding, doesn't seem to be a verbal affair; it pertains to the reality of our thought and its epistemic purport. This remark is certainly not decisive against Khalifa's view, but it does seem to shift the burden of proof to the advocate of thinking that the epistemology of understanding is advanced by decisions about vocabulary.

Lipton and Strevens are explicit that they understand the nature of explanation differently. Lipton stays with the more familiar view on which an explanation is a set of sentences one may utter that accounts for a phenomenon.¹³⁰ Strevens, however, notes that what seems to be crucial for our genuine understanding of a phenomenon is our *ability* to explain it, regardless of whether we possess the relevant verbal resources to do so or not.¹³¹ The discussion in Chapter II about Grimm's view, according to which we only understand what we are able to explain, stays neutral between ways of conceiving explanation exemplified by Lipton and Strevens.

What seems important in the differences between how Lipton and Strevens conceptualize

¹³⁰ Lipton (2009, p. 43) writes: "The explanation is propositional and explicit. It is also conveniently argument shaped, if we take the premise to be the explanation proper and the conclusion a description of the phenomenon that is being explained."

¹³¹ As Grimm (2016) notes, this opens the possibility of non-verbal understanding on the part of pre-verbal children and higher mammals.

explanation is the space opened between our ability to explain phenomena and the verbal finished products that are cited as explanations. In Chapter III, I argued that tacit, non-discursive understanding is a mode in which understanding may occur, and that its occurrence may be experienced as the conscious fluency with which we exploit models of target phenomena or solve problems involving such phenomena. One may deny that tacit understanding exists, and consider what Grimm (2016) calls “an articulacy requirement” on how we understand phenomena. Such a requirement would exclude states of mind by which we tacitly or non-discursively conceive of target phenomena from qualifying as states of understanding precisely on account of their tacitness.

Tentatively, the moral I draw from this dialectic is one that should seem commonsensical, and yet does not seem to have been thoroughly explored in the literature. Namely, that how understanding and the ability to explain what is understood are related to each other depends on how we conceive *both* of explanation and of understanding. On some views (e.g., if explanation is necessarily verbalizable yet tacit understanding is possible), being able to explain what is supposedly understood is no requirement for genuinely understanding it. On other views, being able to explain what is understood *is* required for understanding. This requirement may be imposed by any of two families of conceptions. Either our explanatory abilities need not consist in verbal resources, though they may be manifested with their aid, and so tacit understanding may exercise our explanatory abilities notwithstanding. Or, alternatively, tacit understanding is impossible, and any conception by which we genuinely understand something must be articulable into a verbalized explanation of the phenomena understood. My dissertation stays neutral between these options, but affords framing the issue of how explanation and understanding are related so that differences in how we conceive both terms of the relation may be seen as relevant.

Cognitive phenomenology pertains to what it is like to think, what phenomenal characters one's

occurrent thought experiences assume. Cognitive phenomenology is a hotly debated area of contemporary philosophy of mind; the considerations advanced in Chapter IV don't settle those debates, but offers two suggestions.

Firstly, Chapter IV explored how, if at all, phenomenal and epistemic aspects of understanding might be related to each other. And, in order for the very question debated in that chapter to be raised, episodes in which understanding occurs to us should at least *conceivably* be such that what epistemic norms understanding meets may partly consist in having some conscious experiences, or at least be causally correlated with such experiences. And that would seem to support the claim that some cognitive states may have a phenomenal character, *contra* Carruthers and Veillet (2011), who argue against the possibility of constitutively cognitive phenomenal characters.

Secondly, I argued in Chapter IV that the evidence we have available for how epistemic and phenomenal aspects of understanding are related only supports claims advanced with the modal strength of natural necessity. Such claims are consistent with a variety of underlying metaphysical relationships between phenomenal and epistemic aspects of understanding: constitution, causation, or sheer correlation. So, unless reasons are adduced to think otherwise, Chapter IV also supports the claim that conscious experiences might co-occur with understanding *without* any relation of rational causation or metaphysical grounding obtaining between them.

In this section, I have discussed three philosophical issues that the claims my dissertation makes, if they are true, may bear upon: the value of understanding, the relation between understanding and explanation, and cognitive phenomenology.

4. Open questions

I end by considering a number of questions left open by my dissertation. These may be as many avenues open for further research, which I intend to carry out over the next few years.

One question that needs exploring further is how the conclusions I reach relate to great works in the history of philosophy. Here are only very few of the relevant episodes in the history of philosophy that may need to be revisited: how knowledge, understanding and wisdom are related according to Plato and Aristotle; how insight and revelation may alter our understanding of the world and of ourselves according to Augustine; how meditation may lead to understanding according to Descartes, yet unhinged imagination may hinder it according to Spinoza; the relations between self-understanding, self-knowledge, self-deception and self-transformation for Schelling and Hegel; the inescapability of interpretation whenever the evidence in favor of our understanding is insufficient, controversial or both (arguably, most of the time) for Schleiermacher; the role of insight in scientific discovery and creativity for Whewell, Poincaré and Wertheimer; the relation between acquaintance and grasp for Russell; the phenomenal differences between experiences of meaning, intuitions, and experiences of familiarity in understanding logic according to Husserl; the relationships between pre-understanding, reflective understanding, and leading an authentic life of thought for Heidegger; and many more historical episodes might be relevant to how we may ultimately theorize aspects of understanding.

A second question that needs exploring is how my project of inquiring into how phenomenal and epistemic aspects of understanding relate can be extended to other traditional epistemological notions, such as knowledge and justification. Clearly, we don't think there is something it is like to know, in general. Nor do we think there is something it is like to be justified, in general. However, suppose we

distinguish sources of knowledge, and classify what we know accordingly. There may be nothing it is like to know, in general. But if we consider knowledge by testimony, there may be something it is like to *trust* a close friend sharing a story, or to *defer* to an expert who speaks in her line of expertise. So there is, as far as I can tell, no in-principle barrier to trying to offer an account of how phenomenal and epistemic aspects of knowing relate when we work with suitable taxonomies.

A related open question concerns how the project pursued in this dissertation relates to the program sketched by Alston (2005). For Alston, it is plausible to think that general, unqualified notions of justification are not theoretically fruitful. Rather, Alston thinks, we should focus on the notion of positive epistemic status, viz., that in virtue of which a belief (or a decision) gets to count as justified. Positive epistemic statuses include, in Alston's view: being coherent, being supported by the balance of evidence, etc. I have pursued a strategy similar to Alston's, describing understanding according to how it meets a plurality of epistemic norms, such as being able to *explain* what we understand, the degree to which our understanding of something is *accurate*, or the fact that our understanding should manifest our cognitive *skills* and intellectual virtues (of old or newly acquired, there and then). However, if Alston's line of thought is largely correct, surely remarks similar to his could be made concerning what gets to count as *evidence*, what *coherence* consists in, what the nature of *explanation* is, and so on. The implications this may have for current theories of understanding need to be explored.

I should say a word about the *literatures* on understanding – in the plural. Philosophers of the scientific revolution debated understanding under many guises, and often indistinctly from related kinds of epistemic states of mind: as *scientia*, *scientia intuitiva*, reflective knowledge, virtuous knowledge, *a priori* knowledge, etc. The Enlightenment project, in a sense, is a project in clarifying our ideas, and in better understanding what they are about and how they are related to each other. Several debates in the

philosophy of science concern how scientific understanding and explanation are related. These debates spill over into debates into the cognitive science of science, and the history and philosophy of science. Some debates concern the forms of mathematical understanding and the role of practice-specific mathematical inquiry. There are also debates in the philosophy of mind about the role of conscious experience in understanding, and about what epistemic roles intuitions play. There is a debate in the cognitive science of problem-solving about the role of insight, and how discovering solutions relates to creativity. There is a debate about practical rationality, wisdom, prudence, fluent performance and understanding in virtue ethics. And there is a debate concerning whether tacit understanding underwrites our appreciation of works of art. One important benefit of my dissertation is that it casts its net widely: its remarks can be applied to each of these particular debates. In the years to come I plan to explore these connections.

That said, I end with what is perhaps the most important question this dissertation leaves open, concerning the extent to which generalizations about (intellectual) understanding made in the preceding chapters carry weight in the face of the fact that many of the conceptions by which we understand various phenomena are highly domain-specific. At the end of the day, it may be that what it takes to understand a weather pattern by approaching it with two or three distinct models may be radically different – in a way that cannot be bridged except superficially – from, say, understanding the role that religious motifs play in Chagall’s early paintings. Distinct domains of knowledge may come with distinct forms of understanding that are only superficially similar. That wouldn’t make the generalizations in the preceding chapters *false*; but it would question their *usefulness* in illuminating what we find worth having in the states of mind that qualify as states of understanding. This is a serious challenge, one that I think undermines much of the current literature on understanding – scientific, aesthetic, or otherwise. But

properly addressing it seems the object of an entire research tradition (one whose fruit is unsure) rather than of a single text.

I have sketched a few questions this dissertation leaves open. I end with a word of gratitude, as its work could not be carried out with the support, encouragement, questions and objections of a large number of colleagues and friends. A separate page of acknowledgments is devoted to them.

References

- Alfano, M. (2012) Expanding The Situationist Challenge To Responsibilist Virtue Epistemology. *Philosophical Quarterly* 62, pp. 223-249.
- Alston, W.P. (2005) Epistemic desiderata. Additional Epistemic Desiderata. Reprinted in: *Beyond "Justification:" Dimensions of Epistemic Evaluation* (pp. 11-28, 162-169). Ithaca: Cornell University Press.
- Ansburg, P.I., & Hill, K. (2003) Creative and analytic thinkers differ in their use of attentional resources. *Personality and Individual Differences* 34, pp. 1141–1152.
- Antony, L.M. (2016) Bias: Friend or Foe? Reflections on Saulish Skepticism. In: M. Brownstein & J. Saul (eds.), *Implicit Bias and Philosophy, Vol.1 Metaphysics and Epistemology* (pp. 157-188). Oxford: Oxford University Press.
- Avigad, J. (2008) Understanding Proofs. In: P. Mancosu (ed.) *The Philosophy of Mathematical Practice* (pp. 317-351). Oxford: Oxford University Press.
- Baehr, J. (2006) Character, reliability and virtue epistemology. *Philosophical Quarterly* 56, pp. 193-212.
- Bartlett, G. (2018) Occurrent states. *Canadian Journal of Philosophy* 48, pp. 1-17.
- Baumberger, Chr., Beisbart, Cl., & Brun, G. (2016) What is Understanding? An Overview of Recent Debates in Epistemology and Philosophy of Science. In: St. R. Grimm, Chr. Baumberger, Sabine Ammon (eds.), *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. London: Routledge.
- Bealer, G. (2000) A Theory of the A Priori. *Pacific Philosophical Quarterly* 81, pp. 1-30.

- Bengson, J. (2015) A Noetic Theory of Understanding and Intuition as Sense-Maker. *Inquiry* 58, pp. 633-668.
- Bengson, J. (2015b) The Intellectual Given. *Mind* 124, pp. 707-760.
- Bilalić, Merim, McLeod, P., & Gobet, F. (2010) The Mechanism of the *Einstellung* (Set) Effect: A Pervasive Source of Cognitive Bias. *Current Directions in Psychological Science* 19, pp. 111-115.
- Blackburn, S. (2001) Reason, Virtue, and Knowledge. In: A. Fairweather & L. Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility* (pp. 15-29). Oxford: Oxford University Press.
- Boghossian, P. (2014) What is inference? *Philosophical Studies* 169, pp 1–18.
- Boolos, G. (1971) The Iterative Conception of Set. *Journal of Philosophy* 68, pp. 215-231.
- Bourget, D. (2017) The Role of Consciousness in Grasping and Understanding. *Philosophy and Phenomenological Research* 95, pp. 285-318.
- Brogaard, B. (2013) It's Not What it Seems. A Semantic Account of 'Seems' and Seemings. *Inquiry* 56, pp. 210-239.
- Carruthers, P. & Veillet, B. (2011) The Case Against Cognitive Phenomenology, in T. Bayne & M. Montague (eds.) *Cognitive phenomenology*. Oxford: Oxford University Press.
- Cargile, J. (1970) A Note on Iterated Knowings. *Analysis* 30, pp. 151-155.
- Cargile, J. (1995) Justification and Misleading Defeaters. *Analysis* 55, pp. 216-220.
- Cargile, J. (unpublished) On the Understanding, manuscript.

- Carr, E.H. (1961) Causation in History. In: *What is History?* (pp. 113-143). New York: Random House.
- Carter, J.A., & Pritchard, D. (2016) Cognitive Bias, Scepticism and Understanding. In: St.R. Grimm, C. Baumberger & S. Ammon (eds.), *Explaining Understanding* (pp. 272-292). London: Routledge.
- Chase, V.M., Hertwig, R., & Gigerenzer, G. (1998) Visions of rationality. *Trends in Cognitive Sciences* 2, pp. 206-214.
- Chernev, I. (1955) *1000 Best Short Games of Chess*. Touchstone. Available at:
<http://www.chesshistory.com/winter/extra/movesahead.html>
- Chudnoff, E. (2013) Is Intuition Based On Understanding? *Philosophy and Phenomenological Research* 86, pp. 42-67.
- Cianciolo, A.T., Matthew, C., Sternberg, R.J., & Wagner, R.W. (2006) Tacit Knowledge, Practical Intelligence, and Expertise. In: K.A. Ericsson, N. Charness, P.J. Feltovich, & R.R. Hoffman (eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 613-632). Cambridge: Cambridge University Press.
- Conee, E. (2013) Seeming Evidence. In: C. Tucker (ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism* (pp. 52-70). Oxford: Oxford University Press.
- Cooper N. (1994) The Inaugural Address: Understanding. *Proceedings of the Aristotelian Society*, Suppl. Vol. 68, pp. 1-26.
- Dancy, J. (2006) Reasons, Relevance and Salience: A Response to Hookway. *Philosophical Studies* 130, pp. 71-79.
- de Regt, H.W. (2004) Discussion Note: Making Sense of Understanding. *Philosophy of Science* 71, pp. 98-109.

- de Regt, H.W., & Dieks, D. (2005) A Contextual Approach to Scientific Understanding. *Synthese* 144, pp. 137-170.
- Douven, I., & Meijs, W. (2007) Measuring coherence. *Synthese* 156, pp. 405-425.
- Dreyfus, H.L. (2005) Overcoming the Myth of the Mental: How Philosophers Can Profit from the Phenomenology of Everyday Expertise. *Proceedings and Addresses of the American Philosophical Association* 79, pp. 47–65.
- Elgin, C.Z. (2004) True Enough. *Philosophical Issues* 14, pp. 113–131.
- Ewing, A.C. (1941) Reason and Intuition. “Henriette Hertz Trust” Annual Lecture. British Academy.
- Feldman, R. & Conee E. (2001) Internalism Defended. *American Philosophical Quarterly* 38, pp. 1-18.
- Fennell, J. & Baddeley, R. (2012) Uncertainty Plus Prior Equals Rational Bias: An Intuitive Bayesian Probability Weighting Function. *Psychological Review* 119, pp. 878-887.
- Field, H. (1980/2016) My Strategy for Nominalizing Physics, and its Advantages. In: *Science without Numbers* (pp. 42-47). New York: Oxford University Press.
- Field, H. (1989) Realism, Mathematics, and Modality. In: *Realism, Mathematics, and Modality* (pp. 227-281). New York: Blackwell.
- Franklin, R. L. (1981) Knowledge, belief and understanding. *Philosophical Quarterly* 31, pp. 193-208.
- Fridland, E. (2014) They’ve lost control: reflections on skill. *Synthese* 191, pp. 2729–2750.
- Gertler, B. (2012) Understanding the Internalism-Externalism Debate: What is the Boundary of the Thinker? *Philosophical Perspectives* 26, pp. 51-75.
- Giere, R.N. (2010) Scientific Theorizing. In: *Scientific Perspectivism* (pp. 59-95). Chicago & London:

The University of Chicago Press.

Glymour, C. (1980) Explanations, Tests, Unity and Necessity. *Noûs* 14, pp. 31-50.

Gobet, F., & Charness, N. (2006) Expertise in Chess. In: K.A. Ericsson, N. Charness, P.J. Feltovich, & R.R. Hoffman (eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 523-538). Cambridge: Cambridge University Press.

Goldfarb, W. (1992) Wittgenstein on Understanding. *Midwest Studies in Philosophy* 17, pp. 109-122.

Goldie, P. (2008) Misleading emotions. In: G.Brun, U.Dogluoglu & D.Kuenzle (eds.), *Epistemology and Emotions* (pp. 149-165). Burlington: Ashgate.

Goldman, Al. I. (1997) Science, publicity, and consciousness. *Philosophy of Science* 64, pp. 525-545.

Gopnik, A. (1998) Explanation as Orgasm. *Minds and Machines* 8: 101-118.

Greco, J. (1999) Agent reliabilism. *Philosophical Perspectives* 13, pp. 273-296.

Greco, J. (2014) Episteme: Knowledge and Understanding. In: K. Timpe & C.A.Boyd (eds.), *Virtues and Their Vices* (pp. 285-301). Oxford: Oxford University Press.

Grimm, St. R. (2006) Is Understanding a Species of Knowledge? *British Journal for the Philosophy of Science* 57, pp. 515-535.

Grimm, St. R. (2008) Explanatory inquiry and the need for explanation. *British Journal for the Philosophy of Science* 59, pp. 481-497.

Grimm, St.R. (2010) The Goal of Explanation. *Studies in the History and Philosophy of Science* 41, pp. 337-344.

Grimm, St.R. (2014) Understanding as Knowledge of Causes. In: A. Fairweather (ed.), *Virtue*

Epistemology Naturalized (pp.329-345). Dordrecht: Springer.

Grimm, St. R. (2016) Understanding and Transparency. In: St. R. Grimm, Chr. Baumberger & S. Ammon (eds.), *Explaining Understanding*. New York: Routledge.

Hafner, J. & Mancosu, P. (2005) The Varieties of Mathematical Explanation. In: P. Mancosu (ed.), *Visualization, Explanation and Reasoning Styles in Mathematics* (pp. 215-250). Dordrecht: Springer.

Hills, A. (2015) Understanding Why. *Noûs* 49, pp. 661-688.

Hookway, Chr. (2000) Regulating Inquiry: Virtue, Doubt, and Sentiment. In: G. Axtell (ed.), *Knowledge, Belief, and Character: Readings in Virtue Epistemology* (pp. 149-162). Lanham: Rowman & Littlefield.

Huemer, M. (2007) Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research* 74, pp. 30–55.

Humphreys, P. (2000) Analytic versus synthetic understanding. In: J.H. Fetzer (ed.), *Science, Explanation, and Rationality: Aspects of the Philosophy of Carl G. Hempel* (pp. 267-286). Oxford: Oxford University Press.

James, W. (1890) Chapter XI and XXVI. In: *The Principles of Psychology*. New York: Henry Holt.
Available at: <https://psychclassics.yorku.ca/James/Principles/prin11.htm>

Jorgensen, L.M. (2014) Seventeenth-Century Theories of Consciousness. In: Ed.N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. Available at:
<https://plato.stanford.edu/archives/win2014/entries/consciousness-17th/>

Kahneman, D. (2002) Introduction. Answering an Easier Question. The Illusion of Understanding. In:

Thinking, Fast and Slow. New York: Farrar, Straus and Giroux.

Kelp, Chr. (2015) Understanding phenomena. *Synthese* 192, pp. 3799-3816.

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

Khalifa, K. (2016) Must Understanding be Coherent? In: St. R. Grimm, Chr. Baumberger & S. Ammon (eds.), *Explaining Understanding*. New York: Routledge.

Kim, J. (1994) Explanatory knowledge and metaphysical dependence. *Philosophical Issues* 5, pp. 51-69.

Kozuch, B. & Kriegel, U. (2015) Correlation, Causation, Constitution: On the Interplay between the Science and Philosophy of Consciousness. In: S. M. Miller (ed.), *The Constitution of Phenomenal Consciousness* (pp. 400-417). John Benjamin's.

Kriegel, U. (2013) The Phenomenal Intentionality Research Program. In: *Phenomenal Intentionality*. Oxford: Oxford University Press.

Kupperman, J.J. (1991) Character and Responsibility. In: *Character* (pp. 47-66). New York: Oxford University Press.

Kvanvig, J. (2003) Knowledge and Understanding. In: *The Value of Knowledge and the Pursuit of Understanding* (pp. 185-203). New York: Cambridge University Press.

Laromiguière, P. (1823) Douzième Leçon. In: *Leçons de philosophie* (Vol. I, pp. 278-316). Paris: Bruno-Labbe. Available at: <https://gallica.bnf.fr/.../bpt6k4963j/f1.item.r=entendement>

Lehrer, K. & Paxson, Th., Jr (1969) Knowledge: Undefeated justified true belief. *Journal of Philosophy* 66, pp. 225-237.

Levinson, R. (1994) Experience-Based Creativity. In: T. Dartnall (ed.), *Artificial Intelligence and*

Creativity: An Interdisciplinary Approach (pp. 161-179). Dordrecht: Springer.

Liepnieks, Al. (1961) *The Chess Psychologist: World Champion Tal*. North American Latvian Chess Association. Ann Arbor: University of Michigan Press.

Lipton, P. (2009) Understanding without Explanation. In: H.W. De Regt, S. Leonelli & K.Eigner (eds.), *Scientific Understanding: Philosophical Perspectives* (pp.43-63). University of Pittsburgh Press.

Lynch, M. (2017) Understanding and Coming to Understand. In: St. R. Grimm (ed.), *Making Sense of the World: New Essays on the Philosophy of Understanding*. Oxford: Oxford University Press.

Maddy, P. (2011) *Defending the Axioms*. Oxford: Oxford University Press.

Mancosu, P. (1996) Philosophy of Mathematics and Mathematical Practice in the Early Seventeenth Century. In: *Philosophy of Mathematics and Mathematical Practice in the Seventeenth Century* (pp. 8-33). Oxford: Oxford University Press.

McDowell, J. (1979) Virtue and Reason. *The Monist* 62, pp. 331-350.

Montmarquet, J. (1987) Epistemic Virtue. *Mind* 96, pp. 482-497.

Moore, A.W. (1998/2000) Chapter Eight. In: *Points of View* (pp. 166-194). Oxford: Oxford University Press.

Moravcsik, J.M.E. (1979) Understanding. *Dialectica* 33, pp. 201-216.

Morewedge, C.K., & Kahneman, D. (2010) Associative processes in intuitive judgment. *Trends in Cognitive Sciences* 14, pp. 435-440.

Nakamura, J. & Csikszentmihalyi, M. (2002) The Concept of Flow. In: C.R. Snyder & S.J. Lopez (eds.), *Handbook of positive psychology*. Oxford University Press. Accessed from: <http://eweaver.myweb.usf.edu/2002-Flow.pdf>

- Nanay, B. (2014) An Experiential Account of Creativity. In: E. S. Paul & S.B. Kaufman (eds.), *The Philosophy of Creativity: New Essays* (pp. 17-38). Oxford: Oxford University Press.
- Newman, M.P. (2017) Theoretical Understanding in Science. *British Journal for the Philosophy of Science* 68. Online first.
- Nolan, D. (2014) Hyperintensional Metaphysics. *Philosophical Studies* 171, pp. 149-160.
- Owens, D. (2000) Rationality and Reflection. In: *Reason Without Freedom: The Problem of Epistemic Normativity*. Oxford: Oxford University Press.
- Peacocke, Chr. (1998/2000) Conscious attitudes, attention, and self-knowledge. In: C. Wright, B. Smith & C. Macdonald (eds.), *Knowing Our Own Minds* (pp. 63-97). Oxford: Oxford University Press.
- Polanyi, M. (1958/2015) Skills. In: *Personal Knowledge* (pp. 49-68). Chicago: The University of Chicago Press.
- Pritchard, D. (2005) Cognitive Responsibility and the Epistemic Virtues. In: *Epistemic Luck* (pp. 181-201). Oxford: Oxford University Press.
- Pritchard, D. (2008) Knowing the answer, understanding, and epistemic value. *Grazer Philosophische Studien* 77, pp. 325-339.
- Pritchard, D. (2014) A Modal Account of Luck. *Metaphilosophy* 45, pp. 594-619.
- Railton, P. (2009) Practical competence and fluent agency. In: D. Sobel & St. Wall (eds.), *Reasons for Action* (pp. 81-115). Cambridge: Cambridge University Press.
- Saatsi, J. (2018) Realism and Explanatory Perspectives. In: M. Massimi and C.D. McCoy (eds.), *Understanding Perspectivism: Scientific Challenges and Methodological Prospects*. London: Routledge.

- Saul, J. (2013) Scepticism and Implicit Bias. *Disputatio* 37, pp. 243-264.
- Sosa, E. (1997) Reflective Knowledge in the Best Circles. *Journal of Philosophy* 94, pp. 410-430.
- Strawson, G. (1994) Understanding-experiences. In: *Mental Reality*. Cambridge, Mass.: MIT.
- Strevens, M. (2013) No understanding without explanation. *Studies in History and Philosophy of Science* 44, pp. 510-515.
- Suppes, P. (1960) A comparison of the meaning and uses of models in mathematics and the empirical sciences. *Synthese* 12, pp. 287-301.
- Tetlock, Ph.E., & Belkin, A. (1996) Introduction. In: *Counterfactual Thought Experiments in World Politics: Logical, Methodological, and Psychological Perspectives* (pp.1-38). Princeton, NJ: Princeton University Press.
- Thagard, P. (2002) The passionate scientist: emotion in scientific cognition. In: P. Carruthers, St. Stich & M. Siegal (eds.) *The cognitive basis of science* (pp. 235-250). Cambridge: Cambridge University Press.
- Trout, J.D. (2002) Scientific Explanation And The Sense Of Understanding. *Philosophy of Science* 69, pp. 213-233.
- Tversky, A. & Kahneman, D. (1973) Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* 5, pp. 207-232.
- van Fraassen, B. (1999) How is Scientific Revolution / Conversion Possible? *Proceedings of the American Catholic Philosophical Association* 73, pp. 63-80.
- Watson, G. (1996) Two Faces of Responsibility. *Philosophical Topics* 24, pp. 227-248.
- Wilkenfeld, D.A. (2013) Understanding as representation manipulability. *Synthese* 190, pp. 997-1016.

- Wilson, J. (2014) No work for a theory of Grounding. *Inquiry* 57, pp. 535–579.
- Wilson, T.D., Centerbar, D.B. & Brekke, N. (2002) Mental contamination and the debiasing problem. In: Th. Gilovich, D.Griffin & D.Kahneman (eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 185-200). New York: Cambridge University Press.
- Wittgenstein, L. (1958) *Philosophical Investigations*. G.E.M. Anscombe (transl.). Oxford: Blackwell.
- Woodward, J. (2003) Introduction and Preview. In: *Making Things Happen: A Theory of Causal Explanation* (pp.3-24). Oxford: Oxford University Press.
- Worsnip, A. (2018) The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research* 96, pp. 3-44.
- Ylikoski, P., & Kuorikoski, J. (2010) Dissecting explanatory power. *Philosophical Studies* 148, pp. 201-219.
- Zagzebski, L. (1996) Using moral theory in epistemology. In: *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge* (pp. 1-28). Cambridge: Cambridge University Press.
- Zagzebski, L. (2001) Recovering Understanding. In: M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (pp. 235-249). Oxford: Oxford University Press.