



Routledge Studies in Epistemology

RATIONAL UNDERSTANDING

FROM EXPLANATION TO KNOWLEDGE

Miloud Belkoniene



“Miloud Belkoniene brings forth nothing short of a unified theory of evidence, justification, knowledge and understanding—no small feat indeed!”

Mikael Janvid, *Stockholm University, Sweden*



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Rational Understanding

This book develops a novel account of the connections between justification, understanding, and knowledge. It lays the foundation for a more systematic and interconnected treatment of these central notions in epistemology.

The author's key move is to show first that a specific conception of doxastic justification constitutes our best point of entry into questions pertaining to a subject's ability to secure understanding of reality. Second, that the traditional order of analysis when it comes to the connection between understanding and knowledge should be reversed: knowledge itself is best conceived of in terms of a specific type of understanding.

Rational Understanding will appeal to scholars and advanced students working in epistemology and philosophy of science.

Miloud Belkoniene is Research Associate at the University of Glasgow with research interests in epistemology, philosophy of science, and philosophy of mind. His work on justification and understanding has appeared in journals such as *Synthese*, *Erkenntnis*, and the *British Journal for the Philosophy of Science*.

Routledge Studies in Epistemology

Edited by Kevin McCain, *University of Alabama at Birmingham, USA* and Scott Stapleford, *St. Thomas University, Canada*

Epistemic Autonomy

Edited by Jonathan Matheson and Kirk Lougheed

Epistemic Dilemmas

New Arguments, New Angles

Edited by Kevin McCain, Scott Stapleford and Matthias Steup

Proposition and Doxastic Justification

New Essays on Their Nature and Significance

Edited by Paul Silva Jr. and Luis R.G. Oliveira

Epistemic Instrumentalism Explained

Nathaniel Sharadin

New Perspectives on Epistemic Closure

Edited by Duncan Pritchard and Matthew Jope

Epistemic Care

Vulnerability, Inquiry, and Social Epistemology

Casey Rebecca Johnson

The Epistemology of Modality and Philosophical Methodology

Edited by Anand Jayprakash Vaidya and Duško Prelević

Rational Understanding

From Explanation to Knowledge

Miloud Belkoniene

For more information about this series, please visit: <https://www.routledge.com/Routledge-Studies-in-Epistemology/book-series/RSIE>

Rational Understanding

From Explanation to Knowledge

Miloud Belkoniene



Routledge
Taylor & Francis Group

NEW YORK AND LONDON

First published 2023
by Routledge
605 Third Avenue, New York, NY 10158

and by Routledge
4 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2023 Miloud Belkoniene

The right of Miloud Belkoniene to be identified as author of this work has been asserted in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

The Open Access version of this book, available at www.taylorfrancis.com, has been made available under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 license.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

ISBN: 978-1-032-34877-3 (hbk)

ISBN: 978-1-032-34878-0 (pbk)

ISBN: 978-1-003-32426-3 (ebk)

DOI: 10.4324/9781003324263

Typeset in Sabon
by Deanta Global Publishing Services, Chennai, India

Contents

<i>Preface</i>	viii
Introduction	1
1 On Justification	8
2 The Distance from Truth	29
3 On Understanding	42
4 The Centrality of Truth	59
5 On Knowledge	83
6 The Possession of Truth	107
<i>Index</i>	125

Preface

This book develops a novel account of the connections between justification, understanding, and knowledge. Its guiding idea is that a specific conception of doxastic justification constitutes our best point of entry into questions pertaining to a subject's ability to secure understanding of reality and that addressing those questions can, in turn, shed light on more classical epistemological problems related to knowledge.

Writing a book such as this one is not an individual endeavor as the ideas and arguments presented in it are the result of intense and fruitful philosophical discussions. I am, therefore, extremely grateful for the comments I received from philosophers I had the chance to meet and work with. In particular, I would like to thank Fabian Dorsch (1974–2017), Christoph Kelp, Kevin McCain, Anne Meylan, Gianfranco Soldati, José Zalabardo, the members of the COGITO research group at the University of Glasgow, the members of the ZEGRa research group at the University of Zurich, the members of the EXRE research group at the University of Fribourg, and the members of the GRE research group at Collège de France. Thanks are also due to the Swiss National Science Foundation who provided the funding that made the research presented in this book possible. Finally, I owe special thanks to my *tendre moitié*, Ann-Kathrin, and to my family for their love and support.

Several of the following chapters draw on and systematize previously published work. Chapter 1 partly draws on Belkoniene (2020). Chapter 2 is a significantly revised version of Belkoniene (2021). Chapter 3 is based on Belkoniene (forthcoming a). Chapter 4 is partly based on Belkoniene (forthcoming b).

References

- Belkoniene, M. (forthcoming a). Grasping in Understanding. *British Journal for the Philosophy of Science*.
- Belkoniene, M. (forthcoming b). Reassessing Lucky Understanding. *Episteme*.
- Belkoniene, M. (2021). Why Explanatory Considerations Matter. *Erkenntnis*, 86, 473–491.
- Belkoniene, M. (2020). What Should We Believe about the Future? *Synthese*, 197, 2375–2386.

Introduction

Epistemology can be viewed as an inquiry into the conditions under which knowledge of a mind-independent reality can be secured. While we act and think as if we had knowledge of the world we inhabit, the question as to what it takes, precisely, to secure such an access to reality resurfaces as soon as we adopt a more reflective attitude. For, in acting and thinking we precisely take ourselves to be responding to a mind-independent reality we have access to.

Such an inquiry is usually conceived of as an inquiry into the conditions under which factual knowledge can be gained by a particular subject. When raising questions pertaining to our ability to secure knowledge of reality, we are usually interested in the following question where p stands for a true proposition believed by a subject S :

(Q_K) What does it take for S to know that p is the case?

A good deal of the post-Gettier literature has been devoted to offering a systematic answer to this question and, as underlined by philosophers such as Craig (1990) or Williamson (2000) who adopt a skeptical stance regarding the very possibility of offering such an answer, this has proven to be particularly difficult. Yet, it is worth noting with Greco (2014, p. 285) that a philosopher such as Aristotle would probably not have framed that question in such a way, for the main concern of Aristotle when he talks of episteme—often translated as knowledge—seems to be our ability to account for the world we inhabit. In his view, to possess episteme involves being in a position to account for why p is the case.

It should be clear that the notion of episteme Aristotle operates with differs importantly from the one we tend to operate with when focusing on a question such as Q_K . For there is a clear sense in which a subject can know that p is the case although she ignores why it is the case. Realizing, for instance, that one lacks an account of why Mercury orbits around the Sun is no reason for taking oneself to ignore whether Mercury orbits around the Sun. Aristotle, in fact, seems to be mainly interested in a subject's ability to gain an *understanding* of certain facts; his question being:

2 Introduction

(Q_U) What does it take for S to understand why p is the case?

But why put such an emphasis on Q_U ? As just noted, one does not need an account of why p is the case to be in a position to know that p and it is far from clear that, more generally, securing a cognitive access to reality is dependent on being able to account for it. According to Zagzebski (2009) two main values have, historically, been the target of epistemological inquiry: understanding and certainty. At certain periods, an emphasis was put on questions pertaining to the conditions under which a subject can gain an understanding of reality and, during those periods, the notion of knowledge itself tended to be explicated in relation to the value of understanding. The question as to how an access to reality can be secured tended to be understood relative to the question as to how an account of reality can be provided. At different periods, characterized by a pronounced concern for skeptical problems, the value of certainty was the central one:¹

Skeptical periods have generally been accompanied by the concern for certainty and the process of justifying belief, since justification is what is needed to defend the right to be sure. In contrast, the nonskeptical periods have been mostly concerned with understanding, and the questions accompanying it show little concern for justification but, instead, an interest in the process of explanation, since the ability to explain displays one's understanding.

(Zagzebski, 2009, p. 6)

Attributing to a philosopher such as Aristotle the view that factual knowledge cannot be secured without having an account of the facts one comes to know might thus be rather uncharitable. As suggested by Zagzebski, such philosophers tend to conceive of the main epistemological question as pertaining to a subject's ability to account for reality without thereby denying that there is a sense in which a subject can secure a cognitive access to reality without being in a position to account for it. The reason for placing such a weight on Q_U is, presumably, that if not accompanied by an ability to account for reality such a cognitive access would be dramatically impoverished. While a subject might be able to come to know various facts, it is by accounting for what she knows to be the case that she is able to constitute a unified picture of reality.

This is not to say, however, that there is no priority, in the order of questions, between Q_K and Q_U . For while it is intuitively correct that a subject can secure factual knowledge without having an account of the facts she comes to know, it is not clear that being able to account for p does not require being in a position to know that p in the first place. Suppose with the skeptics that contrary to what we tend to assume, we are in no position to know most of the claims about reality that we take to be true. Is there still a sense, once this supposition is in place, in which we are in a position to provide an account of reality? After all, if the skeptics are correct, although

we may be able to formulate an explanation for p , we are in no position to know whether p requires an explanation in the first place.

This suggests that Q_K actually has priority over Q_U and that the question of the warrant or justification a subject has for endorsing certain claims about the world is as important to philosophers engaged in the project of providing an answer to Q_U as it is to philosophers mostly preoccupied with skeptical problems. Indeed, the notion of justification is taken by several philosophers as being central when it comes to providing an answer to Q_K . As already outlined by Plato in his *Theaetetus*, having knowledge of certain truths involves more than having correct opinions and, although Gettier cases have put substantial pressure on traditional analyses of knowledge, it is plausible that possessing factual knowledge is dependent, in part, on having justification for taking certain propositions to be true.² Skeptical arguments, for instance, do not purport to show that we lack knowledge of a mind-independent reality because the propositions we take to be true are, as a matter of fact, false. They purport to show, instead, that while our beliefs about reality may all be true, we lack justification for rejecting radical hypotheses that, if true, would entail the falsity of most of our beliefs.

If Q_K has priority over Q_U , it is, therefore, reasonable to think that, contrary to what is suggested by Zagzebski, a concern for understanding should be accompanied with a concern for a subject's ability to gain and retain justification for believing certain claims. For if understanding why p involves knowing that p , the question as to what it takes for a subject to understand why p cannot be resolved independently of the question as to what it takes for her to know that p is the case. In fact, one may go as far as arguing that to understand why p simply amounts to knowing why p and that one's ability to account for reality essentially amounts to one's ability to know explanations of phenomena. Of course, given such a view of understanding, Q_K is prior to Q_U in a much stronger sense as, according to this view, our ability to gain an understanding of reality does not raise distinctive epistemological problems: understanding is a species of factual knowledge.³

Such a view of understanding has, however, come under attack in recent literature. Philosophers such as Kvanvig (2003) and Pritchard (2010) argue that factual knowledge is not sufficient to secure an understanding of reality. One important reason they offer is that, for a subject to gain an understanding of why p , the subject needs a particular grasp of the account she has for p ; a grasp that is not required to know that account. Suppose for instance that an account of why water solidifies at a certain temperature is offered to a subject by a physicist who actually understands why this is the case. The account that is offered to the subject states that below a certain temperature, water molecules are sufficiently slow to hook onto each other to form a crystal. Most would admit that the subject to whom this account is provided can come to know why water solidifies at a certain temperature based on the physicist's testimony. Yet, the subject may not grasp the account in such a way as to understand why water solidifies at a certain temperature. She may not see how the movement of water molecules relates

4 Introduction

to the bounds between water molecules or to the solidification of water. This, presumably, does not prevent her from knowing the account that is provided to her but it surely prevents her from understanding why water solidifies at a certain temperature.

Although gaining an understanding of reality may involve more than knowing certain propositions to be true, as just noted, it is *prima facie* plausible that securing factual knowledge is necessary to gain such an understanding. Yet, this claim has also been put into question and there are reasons for thinking that knowledge and understanding in fact have quite different epistemic profiles. In particular, arguments have been offered to the effect that understanding, contrary to knowledge, tolerates certain forms of epistemic luck and that, therefore, knowing why p is the case is not even necessary to secure an understanding of why p .

The question of Q_K 's priority over Q_U is thus more complicated than it initially appears and if the arguments offered against the view that understanding is ultimately a type of factual knowledge are sound, inquiring into the conditions under which a subject knows that p might turn out to consist of a philosophical project that importantly differs from the one characterizing what Zagzebski refers to as non-skeptical periods. The purpose of the present book, however, is to show that these two philosophical projects are in fact closely connected although understanding cannot simply be conceived of as a type of factual knowledge. More precisely, my aim is to show that a certain conception of justification which is plausible in its own right constitutes our best point of entry into questions pertaining to our ability to gain an understanding of reality and that, in turn, questions concerning the conditions under which knowledge can be gained are best addressed by relying on the notion of understanding. As I will argue, both knowledge and understanding, as epistemic phenomena, can be explicated in light of the rational demands imposed on subjects by the constituents of their evidence and those demands have to do with how their evidence can be explained and understood. Here is a brief overview of the issues that will be addressed in this book.

In Chapter 1, I provide a preliminary defense of an explanationist view of evidential support. The view—J-Explanationism—is derived from earlier proposals made by philosophers such as Conee and Feldman (2008) and explicates the conditions under which the evidence possessed by a subject supports believing a particular proposition in terms of how that proposition explanatorily coheres with the evidence. I consider several difficulties raised for such an approach and show that most of them can be overcome. One remaining difficulty, however, concerns the justificatory impact of probabilistic considerations. While some cases support the conclusion that such considerations have weight when it comes to determining whether a subject is justified in believing a particular claim, J-Explanationism suggests that they have none.

Chapter 2 starts, at a more abstract level, with the question as to why one should think that explanatory considerations are connected to the truth

of an explanation. By relying on the notion of truthlikeness introduced by Popper (1963) and, in particular, on recent improvements of Popper's initial construal proposed by Schurz and Weingartner (2010), I show precisely how explanatory considerations are connected to expectations regarding the truth-content of the explanations a subject has for her evidence. This, in turn, allows formulating an updated version of J-Explanationism which shows how probabilistic considerations bear on evidential support and, as a result, is able to deliver the right results in the problematic cases examined in Chapter 1.

If evidential support is closely connected to explanatory considerations, it is reasonable to think that understanding depends on the justification one has, given the evidence one possesses, for holding certain beliefs about the world. Yet, it is also clear that understanding involves more than being justified in believing certain propositions concerning the explanatory story of a phenomenon. In particular, to understand a phenomenon based on an explanation of that phenomenon, a subject must grasp the content of the explanation in the required way. She must, so to speak, see how the explained phenomenon depends, according to the explanation, on the elements cited in that explanation. Chapter 3 investigates the nature of this particular requirement on understanding and, based on earlier proposals made by Newman (2012, 2013, 2017), provides a defense of a view which elucidates this requirement in terms of a specific kind of inferential knowledge; knowledge which is shown to bear important similarities to practical knowledge.

Chapter 4 examines another intuitive requirement on understanding: its accuracy requirement. I argue that the notion of a proposition being central to the understanding of a phenomenon is crucial to the elucidation of the accuracy requirement and offer a detailed account of what it takes for a true proposition to be central in that sense. This account is shown to have plausible consequences and to be, when the context-sensitive nature of understanding is taken into consideration, in accordance with the claim that understanding comes in degrees. Based on the elucidation of the accuracy requirement of understanding and on the considerations put forward in Chapter 3, I then offer a systematic account of explanatory understanding and examine whether knowledge of the propositions that are central to the understanding of a phenomenon is necessary for understanding. I show that propositional knowledge and understanding have, in fact, different epistemic profiles due to their respective sensibility to epistemic luck and that, therefore, knowing the propositions that are central to the understanding of a phenomenon is not required to secure an understanding of that phenomenon.

Chapter 5 investigates the notion of knowledge and offers an elucidation of that notion in terms of a specific type of explanatory understanding. I argue for the view that the notion of knowledge essentially tracks situations in which the evidence possessed by a subject suffices to explain why

6 Introduction

that subject is correct concerning a particular fact and provide a detailed examination of such situations. I then turn to skeptical problems and show how and why the proposed view of knowledge has the resources to address such problems.

Chapter 6 focuses on the notion of evidence which is central to the elucidation of the notions of justification and knowledge put forward in Chapters 1, 2, and 5. I argue that both internal and external facts can constitute the evidence possessed by a subject for believing certain claims about the world and examine the question of the access a subject has to her evidence. Following philosophers such as Pritchard (2011a, 2011b, 2012), I argue that, in epistemically favorable circumstances, a subject has a reflective access to her own evidence but propose a specific construal of that access which, I claim, avoids common difficulties faced by Epistemological Disjunctivism. I then turn to the question of what it takes for something to be part of a subject's evidence and endorse the view that evidence is constituted by the facts—both internal and external—that are manifest to a subject. Finally, I build on my discussion of the notion of evidence to offer a general picture of the connections between justification, understanding, and knowledge. In particular, I show to what extent justification can be conceived of as being directed at gaining an explanatory understanding of the facts that are manifest to a subject and how the proposed elucidation of the notion of knowledge is connected to that dimension of justification.

Notes

- 1 The work of Hellenistic philosophers and modern philosophers such as Descartes or Hume is thought of, by Zagzebski, as being driven, in most part, by a concern for certainty whereas the work of philosophers such as Plato and Aristotle manifests, in her view, a pronounced concern for understanding and the process of explanation.
- 2 This is not to say, of course, that there are no controversies regarding the place the notions of warrant or justification should have in attempts at providing an answer to Q_k . Virtue theoretic accounts of knowledge, for example, can be seen as attempts at providing such an answer without relying on these notions.
- 3 As noted by Grimm (2006) this view of understanding has a venerable tradition within the philosophy of science. See, for instance, Achinstein (1983), Salmon (1989), Woodward (2003), or Lipton (2004). For a recent defense, see Sliwa (2015).

References

- Achinstein, P. (1983). *The Nature of Explanation*. New York: Oxford University Press.
- Conee, E. & Feldman, R. (2008). Evidence. In Q. Smith (Ed.). *Epistemology: New essays*. Oxford: Oxford University Press, 83–104.
- Craig, E. (1990). *Knowledge and the State of Nature*. Oxford: Clarendon.

- Greco, J. (2014). Episteme: Knowledge and Understanding. In K. Timpe & C. A. Boyd (Eds.), *Virtues and Their Vices*. Oxford: Oxford University Press, 285–302.
- Grimm, S. (2006). Is Understanding a Species of Knowledge? *The British Journal for the Philosophy of Science*, 57, 515–535.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lipton, P. (2004). *Inference to the Best Explanation*. New York: Routledge.
- Newman, M. (2012). An Inferential Model of Scientific Understanding. *International Studies in the Philosophy of Science*, 26, 1–26.
- Newman, M. (2013). Refining the Inferential Model of Scientific Understanding. *International Studies in the Philosophy of Science*, 27, 173–197.
- Newman, M. (2017). Theoretical Understanding in Science. *British Journal for the Philosophy of Science*, 68, 571–595.
- Popper, K. R. (1963). *Conjectures and Refutations*. London: Routledge.
- Pritchard, D. (2010). Knowledge and Understanding. In A. Millar, D. H. Pritchard & A. Haddock (Eds.), *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press, 3–91.
- Pritchard, D. (2011a). Epistemological Disjunctivism and the Basis Problem. *Philosophical Issues*, 21, 434–455.
- Pritchard, D. (2011b). Evidentialism, Internalism, Disjunctivism. In T. Dougherty (Ed.), *Evidentialism and Its Discontents*. Oxford: Oxford University Press, 235–253.
- Pritchard, D. (2012). *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Salmon, W. (1989). Four Decades of Scientific Explanation. *Minnesota Studies in the Philosophy of Science*, 13, 3–219.
- Schurz, G. & Weingartner, P. (2010). Zwart and Franssen’s Impossibility Theorem Holds for Possible-world-accounts but not for Consequence-accounts to Verisimilitude. *Synthese*, 172, 415–436.
- Sliwa, P. (2015). Understanding and Knowing. *Proceedings of the Aristotelian Society*, 115, 57–74.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- Zagzebski, L. (2009). *On Epistemology*. Belmont: Wadsworth.

1 On Justification

1.1. *Evidence as What Justifies Belief*

Evidence, one might say, consists of “*what one has to go on* in arriving at a view” (Kelly, 2008, p. 942). While quite broad, this characterization of evidence already presupposes different things. For one, it presupposes that evidence can be possessed by a subject.

Plausibly, the fact that the streets are wet is evidence for the conclusion that it just rained. Likewise, Bob’s footprints on a crime scene can be evidence for the conclusion that Bob committed the crime. But in what sense can that evidence be possessed by a particular subject? Surely, when we say that Bob’s footprints are part of the evidence a subject has we do not mean that the subject literally possesses Bob’s footprints. What we mean is that the subject bears a relation to that evidence in virtue of which it is part of what she has to go on in arriving at a view. We mean that she bears a relation to the fact that Bob’s footprints are on the crime scene that is such that it is part of what she can *use* in arriving at a view—*e.g.*, that Bob is guilty of the crime.

A paradigmatic example of the way evidence can be used by a subject is good reasoning. A subject can use what she observes—that the streets are wet—to arrive, through a particular reasoning, at the conclusion that it just rained. Similarly, a subject who just discovered Bob’s footprints on a crime scene can use that as evidence, in a particular reasoning, to arrive at the conclusion that Bob is guilty of the crime. What makes a particular bit of reasoning “good” is that what is used as a premise in that reasoning bears certain rational relations—*i.e.*, logical, explanatory, or probabilistic relations—to the proposition that constitutes its conclusion. By virtue of bearing such relations to the conclusion of a reasoning, the premises bear on the *truth* of that conclusion.¹

Of course, to say that evidence is what bears on the truth of certain claims is simply another way of expressing the thought that evidence is what supports or justifies holding certain beliefs about the world. Such a view has a long tradition in philosophy. Hume (2000, 10.1), for instance, claims that a “wise man proportions his belief to the evidence” and for Clifford

(2003, p. 518) “it is wrong always, everywhere, and for anyone to believe anything upon insufficient evidence.” If evidence is, in the most basic sense, what bears on the truth of one’s thoughts, it is indeed natural to conceive of the justification one has for holding certain beliefs as depending on the evidence one possesses and, as Williamson puts it, “it is far from obvious that any belief is justified in the truth-directed sense without being justified by evidence” (2000, p. 207).²

It should be noted, however, that several philosophers operate a distinction between two notions of justification pertaining to the attitude of belief.³ In particular, they draw a distinction between propositional and doxastic justification; propositional justification being taken as a property of a subject’s belief’s content and doxastic justification as a property of her belief. The motivation for distinguishing propositional from doxastic justification comes from the observation that while a subject’s evidence can support believing a particular proposition, a subject’s belief in that proposition may fail to be properly based on her evidence. In such a situation, the subject can be said to be propositionally justified in believing that proposition even if her belief in that proposition is not doxastically justified. If such a distinction can be drawn, then the justificatory status of a subject’s belief does not simply depend on the fact that her evidence supports believing its content. It also depends on the fact that her belief is properly based on the evidence she possesses.

In the present book, I will take for granted the view of evidence as what justifies (propositionally) one’s beliefs and focus on certain questions that arise from this view.⁴ In particular, in light of the above considerations, at least three questions deserve further attention:

- (Q₁) What can count as S’s evidence?
- (Q₂) What does it take for something to be part of S’s evidence?
- (Q₃) What does it take for S’s evidence to support believing a particular proposition?

This chapter mainly focuses on the last of these questions as I consider it to be the best point of entry into questions pertaining to a subject’s ability to gain understanding and knowledge of reality; those issues, in turn, have the potential to shed light on the first two questions.

1.2. Evidence as What Is in Need of Explanation

As a first approximation, the evidence possessed by a subject can be said to support believing a proposition *p* whenever *p* fits that evidence. But what does it precisely take for *p* to fit the evidence possessed by a subject? According to explanationist approaches of evidential support defended by philosophers such as Conee and Feldman (2008), McCain (2013, 2014a, 2015, 2017), and Poston (2014), the conditions under which a subject’s

evidence justifies believing that p are to be accounted for in terms of the explanations of that evidence that are available to the subject.

The plausibility of such approaches stems from the fact that we often treat our evidence in such a way. We justify our endorsement of certain claims by relying on explanatory considerations pertaining to the simplicity or the explanatory power of those claims. We show that they ought to be endorsed in light of our evidence because they manage to explain it in a comprehensive and parsimonious way compared to competing claims. For instance, we justify our endorsement of the theory of evolution by pointing out that this theory is best placed to make sense of the evidence we possess. We do not merely treat our findings concerning fossil remains as signs that this theory is true. We treat these findings as being in need of being explained and as justifying endorsing a particular theory of the way living organisms evolve because it represents our best attempt at making sense of these findings. This is corroborated by the important role played by abductive inferences—inferences from a given body of evidence to an explanatory hypothesis—in both scientific and everyday reasoning. The ubiquity of such inferences is outlined by several philosophers⁵ and by experimental studies that tend to show that explanatory considerations have an important impact on the way subjects revise their beliefs.⁶ Thus, evidence does not appear to simply consist of what a subject has to go on in arriving at a view in the sense of being a sign that a particular view is correct. It appears to consist of what a subject has to go on in arriving at a view in the sense of standing in need of being explained by the views a subject might arrive at.

Explanationist approaches of evidential support, as just noted, take the conditions under which a subject's evidence justifies believing that p to be analyzable in terms of the explanations of the evidence that are available to the subject. But, of course, evidential support is not relative to *any* explanation a subject may have for her evidence. In Conee and Feldman's view:

The fundamental epistemic principles are principles of *best explanation* [emphasis added]. Perceptual experiences can contribute toward the justification of propositions about the world when the propositions are part of the best explanation of those experiences that is available to the person. Similarly, the truth of the contents of a memory experience may be part of the best explanation of the experience itself. Thus, the general idea is that a person has a set of experiences, including perceptual experiences, memorial experiences, and so on. What is justified for the person includes propositions that are part of the best explanation of those experiences available to the person.

(2008, pp. 97–98)

Explanations, as suggested by Strevens (2013, p. 510), can broadly be regarded as sets of propositions that instantiate a certain structure; the structure instantiated by a particular explanation depending on the type

of explanation at issue. Presumably, a nomological argument and a causal-mechanistic explanation instantiate different structures but it is by virtue of its respective structure that each set of propositions can purport to explain a phenomenon.⁷ Now, alternative explanations of the same phenomenon can be ranked based on considerations that pertain to their relative explanatory virtues. Some explanations are more powerful than others. They provide a more comprehensive account of what they purport to explain. Some explanations are more simple than others as they provide an account of what they purport to explain that relies on fewer postulates. A particular explanation can, therefore, be deemed, based on such considerations, the best explanation a subject has for her evidence and, according to the view put forward by Conee and Feldman, it is only when p is part of the set of propositions which constitutes the best explanation a subject has for her evidence that p fits her evidence.⁸

Such a view of evidential support, however, appears too restrictive. Consider a case discussed by Goldman (2011, p. 277) of a subject who undergoes a tickle sensation and who attends to that sensation. It should be clear that the subject has justification for believing that she has a tickle sensation. But the reason why she is justified in holding that belief is not that the proposition 'I have a tickle sensation' is part of the best explanation she has for her evidence. A subject who simply undergoes such a sensation and attends to it does not need any explanation of her evidence in order to be justified in believing that she has a tickle sensation. Postulating, as McCain (2014a, pp. 72–73) does, the existence of an intermediary mental state such as the state of attending to the tickle sensation that could be explained by the proposition 'I have a tickle sensation' and that could thereby justify her belief would amount, in my view, to twisting the phenomenon in order to make it fit the theory.⁹ But if a subject does not need an explanation of her evidence containing the proposition 'I have a tickle sensation' to be justified in believing that she has such a sensation, how do such simple cases of introspective beliefs fit within an explanationist approach of evidential support?

The reason why the subject who attends to her tickle sensation does not need an explanation of her evidence containing the proposition 'I have a tickle sensation' to be justified in believing this proposition is, I submit, that when she attends to her sensation the fact that she has a tickle sensation is *manifest* to her. For this reason, the justification the subject has for believing that she has a tickle sensation is substantially different from the justification she would have if the proposition 'I have a tickle sensation' was supported by virtue of being part of the best explanation the subject has for her evidence. As suggested by McDowell (1995) and Neta (2018), when the fact that she has a tickle sensation is manifest to her, the subject is justified in believing that she has such a sensation because that fact itself constitutes her propositional justification for holding that belief. Her justification is immediate. It does not depend on the justification she has for believing another proposition. In contrast, as noted by Fumerton (2018, pp. 337–338), the

justification a subject has for believing a proposition p that is part of the best explanation she has for her evidence is dependent on the justification she has for believing other propositions. In particular, it depends on the justification she has for believing the propositions that constitute her evidence. Indeed, if a subject is not justified in believing the propositions that constitute her evidence and thereby that her evidence stands in need of being explained, her evidence cannot support believing a claim by virtue of being best explained by a particular set of propositions.

Now, how do these considerations help regarding the challenge raised by Goldman for explanationist approaches of evidential support? According to the proposed reading of the case discussed by Goldman, p does not need to be part of the best explanation a subject has for her evidence for her to be justified in believing p and, as a result, principles of best explanation cannot be regarded as being the fundamental epistemic principles. Yet, let me stress that such a reading of the case can itself be motivated by an explanationist approach of evidential support. As just mentioned, the justification a subject has for believing a proposition p that is part of an explanation of her evidence depends on the justification she has for believing that her evidence is in need of being explained. But the justification a subject has for believing that her evidence is in need of being explained cannot itself be explicated in terms of principles of best explanation. Instead, her justification depends on certain facts being manifest to the subject. That is, the reason why a subject is justified in believing that her evidence stands in need of being explained in the first place is plausibly that certain facts are manifest to her. Accordingly, cases in which a subject has justification for believing that p , simply because p is manifest to her, do not undermine the plausibility of what I take to be the central tenet of explanationist approaches of evidential support: that the evidence a subject possesses is adequately conceived of as what is in need of an explanation for that subject. They only illustrate that explanationist approaches of evidential support should be understood as aiming at elucidating the nature of the justification a subject has for believing claims whose content goes beyond what is manifest to her given the (immediate) justification she has for believing what is manifest to her.

In light of these considerations, let me consider the following account which does not require that p be available to a subject as part of an explanation of her evidence for her to be justified in believing p :

J-Explanationism: Whenever p is available to S as part of an explanation for e , S , with e , is justified in believing that p iff p is part of the best explanation S has for e .

This account of evidential support only ranges over situations in which a proposition is available to a subject as part of an explanation for her evidence and, as a result, is more limited in scope than the one put forward by Conee and Feldman (2008). As just argued, I take the propositional

justification a subject has for believing p in situations where p is not available to her as part of an explanation of her evidence to be dependent on p being manifest to that subject, and I shall return to the notion of a fact being manifest to a subject later in this book. For now, let me focus on the types of situation over which J-Explanationism ranges which are, if what I claimed is correct, crucial to understanding the justification a subject can have for believing claims whose content goes beyond what she is immediately justified in believing.

1.3. Explanations and Consistent Beliefs

A first aspect of the account just offered worth highlighting is that e , in J-Explanationism, stands for a subject's *overall* evidence. The reason for this is that unless it does, there is no clear sense in which some piece of evidence e_i could defeat the justification a subject has for believing p . Indeed, if e stands only for a portion of a subject's evidence, by J-Explanationism the subject can be justified in believing that p because p is part of the best explanation she has for that portion of her evidence. But it is not clear how a new piece of evidence e_i could, if acquired by the subject, defeat the justification she has for believing that p , for p would still be part of the best explanation she has for the portion of evidence e . In contrast, if e stands for a subject's overall evidence, a piece of evidence e_i can defeat the justification she has for believing that p if e_i is such that, if added to her overall evidence, p is no longer part of the best explanation she has for that evidence.

One might, however, worry that taking e to stand for a subject's overall evidence makes J-Explanationism overly demanding. Indeed, a subject's overall evidence might include a large amount of propositions and, in normal conditions, subjects may not have access to explanations that could explain all the evidence they possess. Yet, one must first note that a potential explanation of a subject's overall evidence does not need to be such that it would explain every aspect of that evidence. Surely, a potential explanation that would explain every aspect of a subject's evidence would be better, because it would be more explanatorily powerful, than an explanation that would only explain certain aspects of her evidence. But in order for a set of propositions to constitute a potential explanation of a subject's evidence it is not required that that set be such that it would explain every aspect of her evidence.

Second, the conditions under which p is available to a subject as part of an explanation of her evidence should not themselves be construed as being overly demanding. Here I follow McCain (2014a, pp. 66–67) in the way I conceive of those conditions. In particular, I take p to be available to a subject as part of an explanation of why e whenever the subject is disposed to treat p as part of a particular answer to the question “why e ?” upon reflecting on e .¹⁰ According to this construal of the conditions under which p is available to a subject as part of an explanation of e , a subject does not

need to have a particular insight into the way an explanation functions in order for p to be available to her as part of that explanation. Hearing an explanation from someone else without having a particular grasp of the way it works can, in principle, be sufficient for her to be disposed to treat p as part of an explanation of her evidence insofar as her conceptual abilities allow her to access the content of that explanation. For p to be available to a subject as part of an explanation of her evidence, it is therefore not required that the subject actually appreciates the explanatory merits of that explanation in light of her grasp of the way it works.

A second aspect of J-Explanationism that needs to be outlined is that this account need not be understood as being committed to the claim that for p to be part of an explanation of a subject's evidence, p has to contribute to the explanation of that evidence. While it is natural to understand this account in such a way, evidential support is not restricted to propositions that would contribute, if true, to the explanation of a subject's overall evidence. Consider another case discussed by Goldman:

Four Animals on the Deck: I think there are two squirrels on my deck, and I think there are two birds. So I infer that there are (at least) four animals. Presumably, this arithmetic inference is justified. Is it a case of explanatory inference? Surely not. How does there being four animals explain there being two squirrels and two birds? It doesn't. Still, here is a justified belief that some epistemic principle must cover. But that principle, in turn, cannot be grounded in terms of best explanation.

(2011, pp. 277–278)

As Goldman points out, the subject involved in this case appears to be perfectly justified in believing that there are four animals on her deck as a result of her arithmetic inference. But the proposition she infers can hardly be regarded as contributing to the explanation of her evidence. It is simply available to her as a logical consequence of what she is justified in believing—*i.e.*, that there are two squirrels and two birds on her deck and that both squirrels and birds are animals.

While the inference made by the subject involved in this case is clearly not explanatory, I see no reason for thinking that the proposition 'there are four animals on my deck' cannot be part of an explanation of her evidence. The proposition 'there are four animals on my deck' can be regarded as part of the best explanation she has for her evidence despite not contributing itself to the explanation of that evidence. This is because, the justification she has for believing that proposition depends on the fact that the propositions from which she deduces it contribute to a possible explanation of her evidence, which happens to be the best she has.^{11,12} The proposition 'there are four animals on my deck' can, therefore, be regarded as being available to the subject as part of an explanation of her evidence by virtue of being available to her as a logical consequence of propositions that contribute to a possible explanation of her evidence.¹³

Understood in such a way, J-Explanationism encapsulates the idea that if a subject is justified in believing that p given the evidence she possesses, then whenever q is available to her as a logical consequence of p , that subject has sufficient evidence for believing that q . Whenever a proposition q is available to a subject as a logical consequence of p which is itself available to her as part of an explanation of her evidence, q is a further proposition that depends, with respect to its justification, on the quality of the explanation p is part of. Incidentally, the act of explaining itself plausibly involves, on the part of the explainer, a commitment to what logically follows from the propositions which she relies upon to explain a particular phenomenon.¹⁴ If one explains why two squirrels are on the deck by stating that they were attracted by food, one thereby commits to the logical consequences of that statement. If one is not, for instance, ready to endorse the proposition ‘there is food on the deck,’ one can hardly be regarded as explaining why two squirrels are on the deck by means of the statement that they were attracted by food.

As noted by Poston (2014, pp. 100–102), this, however, raises certain difficulties when it comes to the justification a subject has for believing the content of relatively complex explanations. The best explanation a subject has for her evidence can, in principle, consist of a large set of claims which entail a proposition—the conjunction of those claims—that is not probable conditional on the evidence the subject has. While the claims that contribute to a possible explanation of the subject’s evidence might be very probable (yet short of being certain) on that evidence, the probability of their conjunction on that same evidence can be much lower; possibly lower than 0.5. Yet, given the above considerations, if a subject is justified in believing each claim that contributes to a possible explanation of her evidence and their conjunction is available to her as a logical consequence of those claims, she should be justified in believing their conjunction as well. This problem is illustrated by the Preface Paradox put forward by Makinson (1965):

Preface Paradox: Sam has written a book concerning the French Revolution. This book contains a large number of sincerely asserted claims concerning this historical event each of which was carefully examined by Sam. Sam is presently writing the preface of the book and, in that preface, sincerely claims that he is bound to have made some mistakes in his book concerning the French Revolution and that he welcomes any criticisms and comments concerning shortcomings found in his book.

This case presents us with a paradox because Sam appears to be holding beliefs that are inconsistent. Sam sincerely claims, in the preface of the book, that he is bound to have made some mistakes in his book—*i.e.*, that at least one of the claims made in the body of the book is false. Yet, he believes every claim concerning the French Revolution made in the body of

the book. In addition, Sam's evidence appears to support holding a set of inconsistent beliefs. Every claim made in the body of the book concerning the French Revolution has been carefully examined by Sam. Yet, it is very probable on Sam's overall evidence that at least one of those claims is false. But, how can Sam's evidence support believing every claim made in the body of his book and, at the same time, justify believing that at least one of those claims is false?

Given J-Explanationism, Sam cannot, in fact, be justified in believing every claim made in the body of the book concerning the French Revolution and at the same time what is written in the preface. This is due to the fact that the best explanation he has for his overall evidence cannot be logically inconsistent. Such an explanation would not even qualify as a possible explanation of his overall evidence. It is possible that what is written in the preface follows logically from a possible explanation of a portion of Sam's evidence—*e.g.*, from the proposition that every book written by fallible subjects is bound to contain some mistakes. But the proposition expressed in the preface and the propositions expressed in the body of the book cannot all be part of the best explanation Sam has for his *overall* evidence. Either the best explanation of Sam's overall evidence is the one containing the claims made in the body of his book as well as their conjunction or it is the one containing the proposition expressed in the preface of the book.

Because of the explanatory merits of the account of the French Revolution presented in the body of Sam's book, it is reasonable to think that the best explanation Sam has for his overall evidence is actually the one presented in the body of the book. After all, Sam carefully examined each of those claims in light of the historical evidence he managed to gather. But it is also reasonable to think that Sam is justified in believing that the chance that the body of the book contains at least one mistake is high and therefore to interpret what is written in the preface as the expression of a proposition concerning the objective chance of finding mistakes in the body of the book. Note that the proposition 'there is a high chance that the body of my book contains at least one mistake' is not inconsistent with the claims made in the body of the book as their conjunction does not concern the chance of some proposition being true. From p alone one cannot deduce a proposition of the form 'there is a chance c that p ' and therefore p cannot, on its own, be inconsistent with the proposition 'there is a chance c that p ' even if c turns out to be very low.¹⁵ Thus, given J-Explanationism, nothing prevents Sam's evidence from supporting believing both the conjunction of the claims made in the body of his book and that there is a high chance that at least one of the conjuncts is false.

Yet, while this explains how Sam can be justified in believing a set of consistent propositions that takes into account the high probability that his book contains at least one mistake, the intuition that there is something deeply irrational about such a set of beliefs might persist. If a subject's evidence supports the conclusion that the chance of p being false is high, how

can that subject be justified, given that very evidence, in believing that p is true? One way to diffuse this worry is to outline that, by stipulation, Sam's evidence does not point to the falsity of any particular claim made in the body of his book even if it supports believing that the chance of one of those claims being false is high. On the contrary, because of the explanatory merits of the account presented in the book, Sam's evidence points to the truth of each of those particular claims. This should help explain why there is nothing problematic with Sam believing their conjunction while, at the same time, believing that the chance of their conjunction being false is high. As Poston (2014, pp. 101–102) notes, the fact that, given Sam's evidence, the chance of at least one of the claims made in the body of his book being false is high is irrelevant when it comes to the justification he has for believing the account of the French Revolution presented in the book. Sam's justification for endorsing this account depends on its ability to explain the evidence he acquired and, given the explanatory merits of this account, Sam has no reason to believe that any of those particular claims is false.

1.4. *Statistical Evidence and Beliefs about the Future*

The problem illustrated by the Preface Paradox is related to another type of challenge faced by explanationist approaches of evidential support. According to Byerly (2013), such approaches are unable to account for the justification a subject has for holding beliefs concerning the future. This is because propositions about the future cannot explain past and present evidence. To illustrate this problem, Byerly considers the following case:

Golfer: Suppose I'm on the golf course on a sunny, calm day. My putting stroke has been working for me most of the day, and I'm now on the sixteenth green. It's not a long putt—just six feet. I'm fairly confident. I rotate my shoulders, pulling the putter back, and then accelerate through the ball. It rolls toward the cup. The speed looks good. The line looks on. Yes, I believe it's going in!

(2013, p. 235)

The golfer's evidence consists, presumably, of his present observation of the golf ball rolling toward the cup and of his past observations of golf balls rolling toward a cup in similar circumstances. As the claim that the ball *will* go into the cup cannot explain his past and present observations, such a case appears to represent a problem for explanationist approaches of evidential support.

However, as noted in the previous section, J-Explanationism is not committed to the claim that only propositions contributing to the explanation of a subject's evidence can be part of an explanation of her evidence. This account, therefore, has the resources to handle at least some versions of the Golfer case. Suppose, for instance, that the golfer observed that *all* the golf

balls that have rolled toward the cup in circumstances C have gone into the cup and that he is now observing that the golf ball is rolling toward the cup in circumstances C. In this configuration of the case, the best explanation the golfer has for his evidence plausibly includes the generalization ‘all golf balls rolling toward the cup in circumstances C go into the cup’ and the proposition ‘the golf ball is currently rolling toward the cup in circumstances C.’ The generalization, if true, would explain why all the observed golf balls rolling toward the cup in circumstances C have gone into the cup. Now, if the best explanation that the golfer has for his evidence includes these two propositions, the proposition ‘the ball will go into the cup’ can also be available to the golfer as part of the best explanation he has for his evidence and, as a result, the golfer’s past and present evidence can support believing that the ball will go into the cup.

Other configurations of the case are, nevertheless, more difficult to handle. If the golfer observed that *most but not all* of the golf balls that have rolled toward the cup in circumstances C have gone into the cup, the generalization ‘all golf balls rolling toward the cup in circumstances C go into the cup’ can obviously not be part of the best explanation he has for his evidence. To handle such configurations, McCain (2014b, pp. 106–108) considers three possible strategies and I believe that, if properly articulated, the strategy he labels ‘the near neighborhood strategy’ can succeed in answering the challenge raised by the Golfer case.

According to the near neighborhood strategy, while the golfer is not justified in believing the proposition ‘the ball will go into the cup’ in the second configuration of the case, he is justified in believing a proposition that is in the *near neighborhood* of that proposition: that the ball will *probably* go into the cup. This is because, the proposition ‘the ball will probably go into the cup’ instead of the proposition ‘the ball will go into the cup’ is a logical consequence of the best explanation the golfer has for his evidence. Let me suppose, with McCain (2014b, p. 107), that in the second configuration of the case the best explanation the golfer has for his evidence includes the propositions ‘most golf balls rolling toward the cup in circumstances C go into the cup’ and ‘the golf ball is currently rolling toward the cup in circumstances C.’ In light of that supposition, it might be reasonable to think that the proposition ‘the ball will probably go into the cup’ is available to the golfer as a logical consequence of these two propositions and that, therefore, the golfer is justified in believing that the ball will probably go into the cup. Yet, as Byerly and Martin (2014, pp. 778–779) rightly point out, a conjunction of the form ‘most Fs are Gs and x is an F’ does not entail a proposition of the form ‘probably x is a G.’ To see this, consider the following case:

Olympic Runner: Sally is a woman over 35. Suppose most women over 35 are unable to run a 6-min mile. Do these claims entail that it is probable that Sally is unable to run a 6-min mile? They do not. Entailment is supposed to be monotonic. If p entails q , then for any r , $p \wedge r$ entails q .

But, suppose that in addition to being a woman over 35, Sally is a world-class Olympic runner, and that almost all world-class Olympic runners are able to run 6-min miles. If anything, then, it is likely that she can run a 6-min mile.

(Byerly and Martin, 2014, p. 778)

Consequently, the golfer cannot be justified in believing that the ball will probably go into the cup because the propositions ‘most golf balls rolling toward the cup in circumstances C go into the cup’ and ‘the golf ball is rolling toward the cup in circumstances C’ are part of the best explanation he has for his evidence. This, however, does not undermine the viability of the near neighborhood strategy, for the propositions which are relevant for this strategy to succeed are in fact different from the ones just considered. To see why, first note that the term “probably” in the proposition ‘the ball will probably go into the cup’ represents the objective chance of some event to occur. That is, in the context of the near neighborhood strategy, the proposition ‘the ball will probably go into the cup’ is equivalent to the proposition ‘the chance that the ball will go into the cup is higher than 50%’. Accordingly, for the near neighborhood strategy to succeed, what needs to be shown is that the best explanation the golfer has for his evidence includes some propositions that jointly entail that the chance that the ball will go into the cup is higher than 50%.¹⁶

In the second configuration of the case, the golfer’s evidence consists of his observations that most golf balls that have rolled toward the cup in circumstances C have gone into the cup and that the ball is currently rolling toward the cup in circumstances C. It is therefore plausible that the proposition ‘all golf balls rolling toward the cup in circumstances C have a chance of going into the cup higher than 50%’ is part of the best explanation the golfer has for his evidence. This proposition would surely contribute to explain why the golfer observed that the proportion of golf balls that have rolled toward the cup in circumstances C and gone into the cup is higher than the proportion of golf balls that have rolled toward the cup in the same circumstances and not gone into the cup.¹⁷ Now, if the propositions ‘all golf balls rolling toward the cup in circumstances C have a chance of going into the cup higher than 50%’ and ‘the golf ball is currently rolling toward the cup in circumstances C’ are part of the best explanation the golfer has for his evidence, the proposition ‘the chance that the ball will go into the cup is higher than 50%’ is a logical consequence of that explanation. As in the context of the near neighborhood strategy the proposition ‘the chance that the ball will go into the cup is higher than 50%’ is equivalent to the proposition ‘the ball will probably go into the cup,’ the proposition ‘the ball will probably go into the cup’ can be available to the golfer as part of the best explanation he has for his evidence. J-Explanationism is thus able to deliver the result that the golfer is justified in believing that the ball will probably go into the cup.

One may, however, find this result ultimately unsatisfying. After all, given what the golfer observed in the second configuration of the case, it is very probable that the ball will go into the cup. Why think that the golfer is only justified in believing that the ball will *probably* go into the cup? The result delivered by J-Explanationism is in fact in accordance with the distinction drawn by various philosophers between purely statistical evidence and non-statistical evidence. A distinction that is motivated by the apparent inability of statistical evidence to support, on its own, a belief concerning a non-statistical claim.¹⁸ Consider the following situation discussed in the literature concerning legal standards of proof:

Blue Bus: Suppose it is late at night and an individual's car is hit by a bus. This individual cannot identify the bus, but she can establish that it is a blue bus, and she can prove as well that 80 percent of the blue buses in the city are operated by the Blue Bus Company, that 20 percent are operated by the Red Bus Company, and that there are no buses in the vicinity except those operated by one of these two companies. Moreover, each of the other elements of the case – negligence, causation, and, especially, the fact and the extent of the injury – is either stipulated or established to a virtual certainty.

(Schauer, 2003, pp. 81–82)

The evidence available in this case is purely statistical in that it only conveys information about the proportion of blue buses operated in the city by the Blue Bus Company relative to the proportion of blue buses operated in the city by the Red Bus Company. In addition, it seems that the individual whose car got hit by a bus should not, solely on the basis of the evidence available in that case, conclude that the bus that hit her car belongs to the Blue Bus Company. She should only believe that the bus that hit her car *probably* belongs to the Blue Bus Company. This is because, while her evidence makes it probable that the bus that hit her car belongs to the Blue Bus Company, she has no evidence that that bus, in particular, belongs to that company. Her evidence, so to speak, does not point to the particular bus that hit her. As a result, while it may be rational for her to bet that the bus responsible for the accident belongs to the Blue Bus Company, her evidence is not sufficient, on its own, to warrant the conclusion that the bus responsible for the accident belongs to that company. In contrast, if a reliable witness who saw the accident testifies that the bus responsible for the accident belongs to the Blue Bus Company, the testimony of the eyewitness in addition to the statistical evidence present in that case could justify believing that the bus responsible for the accident belongs to the Blue Bus Company. For the testimony could be best explained by the fact that the eyewitness saw that the bus responsible for the accident belongs to the Blue Bus Company. An explanation which contains, by virtue of entailing it, the proposition 'the bus responsible for the accident belongs to the Blue Bus Company.'

In the second configuration of the Golfer case, the evidence possessed by the golfer is also purely statistical with respect to the proposition ‘the ball will go into the cup.’ It only conveys information about the proportion of golf balls that have rolled toward the cup in circumstances C and gone into the cup relative to the proportion of golf balls that have rolled toward the cup in the same circumstances and not gone into the cup. Consequently, it is reasonable to think that the golfer’s evidence does not support the conclusion that the particular ball that is currently rolling toward the cup will go into the cup. It only warrants the conclusion that that ball will probably go into the cup. The distinction between purely statistical evidence and non-statistical evidence also sheds light on the Preface Paradox discussed in the previous section. Recall that I argued that Sam is justified in believing the conjunction of all the claims made in the body of his book. One could nevertheless reason as follows: as the probability of the proposition ‘at least one of the claims made in the body of the book is mistaken’ is higher on Sam’s overall evidence than the probability of the conjunction of all the claims made in the body of the book, if Sam is justified in believing the conjunction of all the claims made in the body of his book, then he should also be justified in believing that at least one of those claims is mistaken. More generally, if a subject is justified in believing that p and q ’s probability on that subject’s evidence is higher than p ’s probability, it is intuitively correct that that subject is also justified in believing that q . Yet, the inability of purely statistical evidence to justify, on its own, believing a non-statistical claim shows that such a reasoning is flawed. In the Preface Paradox, Sam’s evidence is purely statistical with respect to the proposition ‘at least one of the claims made in the body of the book is mistaken.’ For this reason, even if that proposition’s probability is higher on his evidence than the probability of the conjunction of all the claims made in the body of his book, his evidence does not support believing that the body of his book contains at least one mistaken claim. Sam’s evidence supports, at most, believing that this is probably the case.

Generally, if a subject’s evidence can be adequately conceived of as what is in need of an explanation for that subject, one should not expect purely statistical evidence to support, on its own, believing non-statistical claims. Such evidence is to be explained in terms of statistical claims that concern the chance of some event to occur. This is not to say, of course, that all the evidence that a subject can possess in support of claims about the future is purely statistical. If the best explanation a subject has for her evidence entails a non-probabilistic claim about the future, that subject’s evidence supports believing that an event will occur. Moreover, note that this condition can be satisfied even if the best explanation a subject has for her evidence does not include a universally quantified proposition as in the first configuration of the Golfer case. To see why, suppose for instance that an individual is waiting for a shop to open and that someone passing by tells her that the shop will open at 9 a.m. If the individual has no particular

reason to suspect that that person is lying to her or that she is not in a position to know that the shop will open at 9 a.m., it seems that she is justified in believing that the shop will open at 9 a.m. on the basis of that person's testimony. The best explanation she has for that person's testimony is that the person who told her that the shop will open at 9 a.m. knows that it will open at 9 a.m. and as this proposition entails the proposition 'the shop will open at 9 a.m.', her evidence can, in principle, support believing that the shop will open at 9 a.m.

1.5. *Justification and Improbable Beliefs*

The discussion of the Preface Paradox and the Golfer case could be taken to support the conclusion that probabilistic considerations have no central role when it comes to determining whether a subject has justification for believing a proposition. Regarding the Preface Paradox, I argued that Sam's evidence can support believing the conjunction of the propositions expressed in the body of his book even if that conjunction's probability, conditional on Sam's evidence, is low. Regarding the Golfer case, I argued that the golfer's evidence, in the second configuration of that case, does not support believing that the ball will go into the cup even if that proposition's probability conditional on his past and present evidence is high. This naturally suggests that probabilistic considerations have no justificatory import. Yet, such a conclusion deserves a more careful examination. Consider the following situation:

Investigation: Imagine that Sally is the lead detective on an investigation of a burglary. She typically uses an eight-step investigative procedure for crimes of this sort and this procedure involves gathering and analyzing multiple kinds of evidence—physical evidences, forensic evidences, testimonial evidences, psychological evidences, circumstantial evidences, and so on. Sally is now mid-way through her investigative procedure, having completed four of the eight steps. She has gathered and analyzed the appropriate evidence for these four steps, but has not yet gathered or analyzed evidence that may or may not arise during the final four steps. The list of suspects with which Sally began has been narrowed, and there is one very promising suspect in particular named Jeremy. In fact, the claim 'Jeremy committed the burglary' (call this the Jeremy hypothesis) is the best explanation available to Sally for all of the evidence she currently has obtained through the first four steps. There are multiple witnesses locating someone who fits Jeremy's description at the scene of the crime at the time at which it was committed. Some drug paraphernalia like that which Jeremy commonly uses to feed his drug habit was found at the scene of the crime. Jeremy seems to display a sense of satisfaction or gladness about the robbery. His bank account reflects a deposit shortly after the incident. Other current

suspects, while not ruled out, do not fit the evidence Sally currently has anywhere nearly as well as Jeremy does. The Jeremy hypothesis is the best available explanation for the evidence Sally currently has and it is a very good explanation of that evidence.

(Byerly and Martin, 2014, p. 783)

Byerly and Martin add that it often happened to Sally that, after completing the last steps of her investigation procedure, a new suspect emerged that better fitted the evidence she had gathered. Given this additional information, they conclude, rightly it seems, that mid-way through her investigative procedure, Sally is not justified in believing that Jeremy committed the burglary, although the proposition ‘Jeremy committed the burglary’ is part of what could best explain the evidence she has.

How is such a case supposed to be handled if one endorses a view such as J-Explanationism? According to McCain (2015, 2017, 2018) the Jeremy hypothesis is not the best explanation Sally has for her overall evidence mid-way through her investigative procedure. While it might be the best one she has for a portion of her evidence, it is not the best explanation she has for her overall evidence. Mid-way through her investigative procedure, a more general explanation containing the proposition ‘some, as yet unknown, suspect committed the burglary’ explains Sally’s overall evidence better than the Jeremy hypothesis. This is because, while a portion of Sally’s evidence is best explained by the Jeremy hypothesis, another portion consists of strong inductive evidence for the claim that the best explanation Sally has for the evidence she has acquired so far is not the correct one.

This line of answer to the challenge raised by the case Byerly and Martin offer is, however, problematic. First, as Byerly and Martin (2016) and myself (2017) point out, it is not clear that an explanation containing the proposition ‘some, as yet unknown, suspect committed the burglary’ is better *qua* explanation than the Jeremy hypothesis once Sally’s overall evidence is considered. The fact that some, as yet unknown, suspect committed the burglary could not explain Sally’s past experience of investigations. In addition, it is a poor explanation of the evidence Sally gathered during the first four steps of her investigative procedure.

Second, in light of the considerations put forward in the previous sections, one should not expect Sally’s past experience of investigations to support the conclusion, mid-way through the investigative procedure, that the Jeremy hypothesis is not the correct explanation of the evidence she has gathered. In fact, there seems to be no substantial difference between Sally’s situation and Sam’s situation in the Preface Paradox. Sally’s overall evidence appears to support believing at most that the Jeremy hypothesis is *probably* not the correct explanation of the evidence she has gathered so far. As a matter of fact, if Sam’s evidence can support believing at the same time the conjunction of all the claims made in the body of his book and that at least one of those claims is probably mistaken, there is no principled reason for

denying that Sally's evidence, mid-way through her investigative procedure, can support believing that Jeremy committed the burglary. Yet, when it comes to the Investigation case, this result appears unacceptable. Intuitively, Sally ought to complete her investigative procedure before arriving at any definitive conclusion regarding the culprit.^{19,20}

1.6. *Concluding Remarks*

According to J-Explanationism, the view of evidential support that has been examined in the present chapter, a subject's evidence does not simply consist of what she has to go on in arriving at a view in the sense of being a sign that a particular view is correct. A subject's evidence consists of what is standing in need of being explained for that subject and explanatory considerations are therefore crucial when it comes to determining whether a subject is justified in believing certain claims. It does not follow, however, that a subject's evidence can support believing a proposition only insofar as that proposition is available to her as part of an explanation of her evidence. For, as I argued, J-Explanationism is best understood as aiming at elucidating the nature of the justification a subject has for believing claims whose content goes beyond what is manifest to her given the (immediate) justification she has for believing what is manifest to her.

Several challenges raised against explanationist approaches of evidential support have been examined and the discussion that followed showed that the results delivered by J-Explanationism in cases taken to be problematic for such approaches are in line with an independently plausible distinction between purely statistical evidence and non-statistical evidence. This is not to say, however, that the results delivered by J-Explanationism are all intuitively acceptable. Cases such as the Investigation suggest that J-Explanationism, in its present form, fails to capture something important concerning evidential support and, as a result, should not be endorsed without further examination.

Notes

- 1 As Williamson (2000, pp. 194–200) emphasizes, only propositions or facts bear rational relations to other propositions and, in what follows, I will work under the assumption that evidence is best conceived of as being propositional.
- 2 See also Conee and Feldman (1985, 2004, 2011) for a notable defense of the view that facts about the justificatory status of a subject's beliefs supervene on facts about that subject's evidence.
- 3 See for instance Firth (1978), Conee and Feldman (1985), and Pryor (2004).
- 4 While this view constitutes a starting point rather than something I wish to directly argue for, the considerations put forward in this book speak, I believe, in favor of the conclusion that adopting such a perspective is fitted when it comes to elucidating philosophically important notions such as doxastic justification, understanding, and knowledge.
- 5 Regarding the role of abductive inferences in scientific contexts, see McMullin (1992) and Lipton (2004). Regarding the role of such inferences in everyday contexts see Harman (1965) and Adler (1994).

- 6 See Lombrozo and Carey (2006), Douven and Verbrugge (2010), and Lombrozo and Gwynne (2014).
- 7 In the present context, the question as to whether there are really distinct types of explanations will be left open.
- 8 Note that the best explanation a subject has for her evidence consists of the best explanation among the potential explanations of her evidence where H is a potential explanation of some evidence whenever H *would* explain that evidence. As a result, the best explanation a subject has for her evidence does not necessarily consist of the set of true propositions that *actually* explains that evidence—*i.e.*, a correct explanation of the evidence.
- 9 Conee (2020) actually takes this particular problem to be fatal to the view he previously defended.
- 10 Proponents of explanationist approaches of evidential support such as Conee and Feldman (2008) or McCain (2014a) take such approaches to be closely related to an internalist conception of justification. It is important to note, however, that a subject's ability to reflect on *e* need not presuppose an ability to gain access to her evidence on the basis of reflection alone. Even if the subject's evidence is not reflectively accessible to her, she can reflect on that evidence once she has gained an empirical access to it. In Chapter 6, I will return to the question of the type of access a subject has to whatever constitutes her evidence. For now, it is sufficient to note that J-Explanationism need not be understood as being committed to the claim that for a subject to be justified in believing *p* given some evidence *e*, *e* needs to be reflectively accessible to her.
- 11 See Lehrer (1974) for similar counter-examples to explanationist views of evidential support.
- 12 See Poston (2014) who makes similar observations regarding his own account of evidential support and McCain (2013, 2014a) who extends his account of evidential support to relations of logical consequences in light of such considerations.
- 13 This allows J-Explanationism to avoid the Expertise Objection raised by Conee (2020, p. 77). This objection is based on the thought that expert explanations often don't contain the commonsense propositions that we take ourselves to be justified in believing although they are better, *qua* explanations, than the ones containing those propositions. Conee considers a case in which a subject has two explanations for why she feels pain in her wrist. One explanation, which is correct, states that a splash of hot water scalded her wrist. The other, also correct, explains the pain at a cellular level. In Conee's view, because the second explanation is better than the first, an explanationist view of evidential support cannot account for the fact that the subject is justified in believing that a splash of hot water scalded her wrist. Yet, it is not clear that the two explanations can really be regarded as rival explanations. This is because it is plausible that the second explanation also contains the commonsense proposition by way of entailing it.
- 14 See Dellsén (2018, 2019) for a defense of that claim.
- 15 Let me insist on the fact that I am considering here a proposition about the objective chance of *p* being true and not a proposition that concerns *p*'s evidential probability.
- 16 Let me emphasize that for that strategy to succeed, it does not have to be the case that the chance that the ball will go into the cup actually is higher than 50%. It is sufficient that the golfer's evidence supports believing that that chance is higher than 50%.
- 17 One might object that an explanation containing the proposition 'most golf balls rolling toward the cup in circumstances C have a chance of going into the cup (significantly) higher than 50%' is an equally good explanation of the golfer's observations. After all, this proposition could also contribute to explain

why the golfer observed that most golf balls that have rolled toward the cup in circumstances C have gone into the cup. Yet, note that this constitutes an objection only if the proposition ‘most golf balls rolling toward the cup in circumstances C have a chance of going into the cup (significantly) higher than 50%’ entails that a minority of the golf balls rolling toward the cup in circumstances C have a chance of going into the cup lower than 50%. Otherwise, this proposition is equivalent to the proposition ‘all golf balls rolling toward the cup in circumstances C have a chance of going into the cup higher than 50%’. It is, however, doubtful that an explanation which entails that a minority of the golf balls rolling toward the cup in circumstances C have a chance of going into the cup lower than 50% is an equally good explanation of the golfer’s evidence. If the golfer simply observed that most golf balls that have rolled toward the cup in circumstances C have gone into the cup, there is no reason to suppose that a minority of golf balls rolling toward the cup in those circumstances have a chance of going into the cup lower than 50%. An explanation stating that all golf balls rolling toward the cup in those circumstances have a chance of going into the cup higher than 50% is by far the simplest explanation of the golfer’s observations.

- 18 While many philosophers support drawing such a distinction, there is no agreement among them concerning the features that make some evidence non-statistical as opposed to statistical. For possible accounts of non-statistical evidence see: Cohen (1977), Thomson (1986), Dant (1989), Enoch, Spectre, and Fisher (2012), Smith (2016, 2018), and Blome-Tillmann (2017). For explanation-based accounts of non-statistical evidence see Pardo and Allen (2008) and Belkoniene (2019).
- 19 McCain (2018, p. 67) considers another strategy to handle cases such as the one offered by Byerly and Martin which amounts to claiming that while the Jeremy hypothesis might be the best explanation Sally has for her overall evidence, that explanation is not good enough for Sally’s evidence to support believing that Jeremy committed the burglary. While this strategy is available to proponents of explanationist approaches of evidential support, it raises the question as to why the Jeremy hypothesis is not good enough and the only answer here seems to be that it is not probable enough given Sally’s overall evidence. Consequently, this strategy does not spare explanationists from providing a precise account of the justificatory import of probabilistic considerations and of their precise connection with explanatory considerations.
- 20 See Appley and Stoutenburg (2017) for a similar objection against explanationist views of evidential support.

References

- Adler, J. (1994). Testimony, Trust, Knowing. *Journal of Philosophy*, 91, 264–275.
- Appley, B. & Stoutenburg, G. (2017). Two New Objections to Explanationism. *Synthese*, 194, 3069–3084.
- Belkoniene, M. (2017). What Are Explanatory Virtues Indicative Of? *Logos and Episteme*, 8, 179–193.
- Belkoniene, M. (2019). Normalcy, Understanding and the Problem of Statistical Evidence. *Theoria*, 85, 202–218.
- Blome-Tillmann, M. (2017). ‘More Likely Than Not’ Knowledge First and the Role of Bare Statistical Evidence in Courts of Law. In A. Carter, E. Gordon & B. Jarvis (Eds.). *Knowledge First*. Oxford: Oxford University Press, 278–292.
- Byerly, R. (2013). Explanationism and Justified Beliefs about the Future. *Erkenntnis*, 78, 229–243.

- Byerly, R. & Martin, K. (2014). Problems for Explanationism on Both Sides. *Erkenntnis*, 80, 773–791.
- Byerly, R. & Martin, K. (2016). Explanationism, Super-Explanationism, Ecclectic Explanationism: Persistent Problems on Both Sides. *Logos & Episteme*, 7, 201–213.
- Clifford, W. (2003). The Ethics of Belief. In L. Pojman (Ed.). *The Theory of Knowledge*. Belmont: Wadsworth, 515–518.
- Cohen, L. J. (1977). *The Probable and the Provable*. Oxford: Oxford University Press.
- Conee, E. (2020). Evidential Support and Best Explanations. *Philosophical Issues*, 30, 71–85.
- Conee, E. & Feldman, R. (1985). Evidentialism. *Philosophical Studies*, 48, 15–34.
- Conee, E. & Feldman, R. (2004). *Evidentialism*. New York: Oxford University Press.
- Conee, E. & Feldman, R. (2008). Evidence. In Q. Smith (Ed.). *Epistemology: New Essays*. Oxford: Oxford University Press, 83–104.
- Conee, E. & Feldman, R. (2011). Replies. In T. Dougherty (Ed.). *Evidentialism and Its Discontents*. New York: Oxford University Press, 428–501.
- Dant, M. (1989). Gambling on the Truth: The Use of Purely Statistical Evidence as a Basis for Civil Liability. 22 *Columbia Journal of Law and Social Problems*, 31, 31–70.
- Dellsén, F. (2018). Deductive Cogency, Understanding, and Acceptance. *Synthese*, 195, 3121–3141.
- Dellsén, F. (2019). Rational Understanding: Toward a Probabilistic Epistemology of Acceptability. *Synthese*, 198, 2475–2494.
- Douven, I. & Verbrugge, S. (2010). The Adams Family. *Cognition*, 117, 302–318.
- Enoch, D., Spectre, L. & Fisher, T. (2012). Statistical Evidence, Sensitivity, and the Legal Value of Knowledge. *Philosophy and Public Affairs*, 40, 197–224.
- Firth, R. (1978). Are Epistemic Concepts Reducible to Ethical Concepts? In A. Goldman & J. Kim (Eds.). *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*. Dordrecht: Reidel, 215–229.
- Fumerton, R. (2018). Evidentialism and Explanatory Fit. In K. McCain (Ed.). *Believing in Accordance with the Evidence*. Cham: Springer, 329–343.
- Goldman, A. (2011). Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package. In T. Dougherty (Ed.). *Evidentialism and Its Discontents*. New York: Oxford University Press, 254–280.
- Harman, G. (1965). The Inference to the Best Explanation. *Philosophical Review*, 74, 88–95.
- Hume, D. (2000). *An Enquiry Concerning Human Understanding* (T. Beauchamp, Ed.). Oxford: Clarendon Press.
- Kelly, T. (2008). Evidence: Fundamental Concepts and the Phenomenal Conception. *Philosophy Compass*, 3, 933–955.
- Lehrer, K. (1974). *Knowledge*. Oxford: Oxford University Press.
- Lipton, P. (2004). *Inference to the Best Explanation*. London: Routledge.
- Lombrozo, T. & Carey, S. (2006). Functional Explanation and the Function of Explanation. *Cognition*, 99, 167–204.
- Lombrozo, T. & Gwynne, N. (2014). Explanation and Inference: Mechanical and Functional Explanations Guide Property Generalization. *Frontiers in Human Neuroscience*, 8.

- Makinson, D. (1965). The Paradox of the Preface. *Analysis*, 25, 205–207.
- McCain, K. (2013). Explanationist Evidentialism. *Episteme*, 10, 299–315.
- McCain, K. (2014a). *Evidentialism and Epistemic Justification*. New York: Routledge.
- McCain, K. (2014b). Evidentialism, Explanationism, and Beliefs about the Future. *Erkenntnis*, 79, 99–109.
- McCain, K. (2015). Explanationism: Defended On All Sides. *Logos & Episteme*, 6, 61–73.
- McCain, K. (2017). Undaunted Explanationism. *Logos & Episteme*, 8, 117–127.
- McCain, K. (2018). Explanatory Virtues are Indicative of Truth. *Logos & Episteme*, 9, 63–73.
- McDowell, J. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research*, 55, 877–893.
- McMullin, E. (1992). The Inference that Makes Science. *Zygon*, 48, 143–191.
- Neta, R. (2018). Your Evidence Is the Set of Facts That Are Manifest to You. In V. Mitova (Ed.). *The Factive Turn in Epistemology*. Cambridge: Cambridge University Press, 32–49.
- Pardo, M. S. & Allen R. J. (2008). Juridical Proof and the Best Explanation. *Law and Philosophy*, 27, 223–268.
- Poston, J. (2014). *Reason and Explanation*. New York: Palgrave Macmillan.
- Pryor, J. (2004). What's Wrong with Moore's Argument? *Philosophical Issues*, 14, 349–378.
- Schauer, F. (2003). *Profiles, Probabilities, and Stereotyping*. Cambridge: Belknap Press.
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- Smith, M. (2018). When Does Evidence Suffice for Conviction? *Mind*, 127, 1193–1218.
- Strevens, M. (2013). No Understanding Without Explanation. *Studies in History and Philosophy of Science Part A*, 44, 510–515.
- Thomson, J. J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems*, 49, 199–219.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

2 The Distance from Truth

2.1. *Two Hypotheses*

The Investigation case raised an important question concerning the justificatory import of probabilistic considerations. Mid-way through her investigative procedure, Sally does not appear to be justified in believing that Jeremy committed the burglary even if that proposition is available to her as part of the best explanation she has for the evidence she gathered. The reason, it seems, is that while being part of the best explanation of Sally's overall evidence, that proposition is not probable conditional on her evidence. But, how can this explain why Sally is not justified in believing that Jeremy committed the burglary if, in the Preface Paradox, the fact that the conjunction of all the claims made by Sam in the body of his book is not probable on his evidence does not prevent him from being justified in believing that conjunction?

To resolve this apparent tension, it is necessary to take a step back and examine the connection between explanatory considerations and the truth of the explanations a subject has for her evidence in more detail. If one abstracts from the results delivered by J-Explanationism in the cases considered in the previous chapter and brackets this account of evidential support for now, why should one think that explanatory considerations are closely connected to the truth of an explanation's content?

Explanatory considerations, as already noted, are related to the virtues of an explanation. They are considerations of the sort " H_1 would explain why e better than H_2 would, if true." As the virtues of an explanation are properties that are possessed in relation to a body of evidence that stands in need of an explanation, it is reasonable to think that explanatory considerations can shed light on what a body of evidence is evidence for. Yet, such a conclusion needs to be accounted for. What is precisely the relation between the virtues of some explanation and the truth of its content-parts? A first possible answer is that explanatory considerations are a good yet fallible guide to estimate the probability of some claim on the available evidence. In other words, explanatory considerations can shed light on what a body

of evidence is evidence for insofar as they are a good guide to estimate the probability of some hypothesis on that body of evidence.¹

Yet, in light of the considerations put forward in the previous chapter, I believe that such an answer should be rejected. First, if explanatory considerations are merely a good guide to estimate the probability of some claim on the available evidence, it is not clear why the available evidence should be conceived of as what is in need of an explanation for a subject. Evidential support, according to this hypothesis, is fundamentally a matter of probabilistic considerations as philosophers such as Foley (1992, 2009) would have it. Second, explanatory considerations and probabilistic considerations can come apart. As outlined in the previous chapter, a claim can be part of the best explanation a subject has for her evidence without being probable conditional on that evidence. Inversely, a claim can be probable on the available evidence without being part of the best explanation one has for that evidence. Consider for instance a subject who just bought a lottery ticket. If her evidence includes the proposition that the lottery is fair and involves 10,000 tickets, the claim ‘I am holding a losing ticket’ is very probable on her evidence. However, this proposition is not part of an explanation of her evidence.

Here, I intend to develop an alternative hypothesis—the *truthlikeness hypothesis*—that, I will argue, is both plausible and able to shed light on the account of evidential support put forward in the previous chapter. To the question “what is the relation between the virtues of some explanation and the truth of its content-parts?” let me offer the following answer: the explanatory virtues of an explanation H of e reflect H ’s relative closeness to the truth about e . This in the sense that the best explanation a subject has for e can be expected, given e , to be closer to the whole truth than any other explanation of e available to the subject. This hypothesis differs from the one just considered in that it relies on the notion of closeness to the whole truth or truthlikeness as this notion is standardly referred to. The whole truth can be conceived of as the most informative true theory expressible in a language about a given domain. Consequently, it provides, as a theory, the best possible answer to the question “why is e the case?” Now, if H_1 —a potential explanation of e —is better, *qua* explanation, than H_2 , it is *prima facie* plausible that H_1 ’s explanatory virtues reflect its relative closeness to the theory that provides the best possible answer to the question “why is e the case?”—*i.e.*, its relative closeness to the whole truth. This is because H_1 being better, *qua* explanation, than H_2 just means that H_1 would provide a better answer to the question “why is e the case?” than H_2 would, if true.²

Before examining the notion of truthlikeness in more detail, let me offer some clarifications. First, the truthlikeness hypothesis concerns the expected truthlikeness of an explanation and not its actual truthlikeness. It does not state that the best explanation a subject has for e is actually more truthlike than the other explanations of e available to her. Instead, according to this hypothesis, the best explanation a subject has for e can be *expected*, given

e , to be more truthlike than the other explanations of e available to her. The notion of expectation I am relying on here is an epistemic one and can be understood in terms of propositionally justified belief. Accordingly, the claim that H can be expected, given e , to be more truthlike than any other explanations of e available to a subject, amounts to the claim that e supports believing that H is more truthlike than any other explanations of e available to that subject.³ Second, the truthlikeness hypothesis concerns the expected *relative* truthlikeness of an explanation and not its expected absolute degree of truthlikeness. The whole truth provides the best possible answer to the question “why is e the case?” But the information it contains goes far beyond facts pertaining to the explanation of e . For this reason, the best explanation a subject has for e might be expected to share only a little amount of content with the whole truth. That is, while the best explanation a subject has for e can, given e , be expected to be more truthlike than the other explanations of e available to her, it might, at the same time, be expected to be far from the whole truth. Similar considerations apply to the notion of probability. The fact that the best explanation a subject has for e can be expected to be more truthlike than the other explanations available to her does not entail that that explanation is probable conditional on e . The notions of probability and truthlikeness differ quite importantly as it will become clear.

2.2. Truthlikeness, Theories, and Relevant Consequences

Popper (1963) was the first to offer a precise rendering of the notion of truthlikeness.⁴ To that end, he distinguishes the truth-content of a theory from its falsity-content. A theory H 's content, written $C(H)$, is taken by Popper to consist of the set of propositions logically entailed by H and its truth-content $C(H)_T$, a subset of $C(H)$, is defined as the set of H 's true logical consequences. For its part, H 's falsity-content $C(H)_F$, also a subset of $C(H)$, is defined as the set of H 's false logical consequences. Given the distinction between these two subsets of $C(H)$, Popper is able to articulate the idea that if H is true—that is, if $C(H)_F = \{\phi\}$ —something can still be said, given H 's truth-content $C(H)_T$, about that theory's closeness to the whole truth. Likewise, if H is false—if $C(H)_F \neq \{\phi\}$ —something can still be said about that theory's closeness to the whole truth given both its truth-content $C(H)_T$ and its falsity-content $C(H)_F$. Whether H is true or false, $C(H)_T$ and $C(H)_F$ allow determining whether H is more truthlike than another theory, provided that the two theories are comparable.

Popper (1963, p. 233) formulates two disjunctively sufficient conditions for a theory H_1 to be more truthlike than another theory H_2 . For H_1 to be more truthlike than H_2 , the truth-content of H_1 (but not its falsity-content) must be a proper subset of H_2 's truth-content or the falsity-content of H_1 (but not its truth-content) must be a proper subset of H_2 's falsity-content. Popper's definition can be formally represented as follows, where H_1 being more truthlike than H_2 is written $H_1 >_T H_2$:

P-Truthlikeness: $H_1 >_T H_2$ iff:

- (i) $C(H_1)_T \subset C(H_2)_T$ and $C(H_1)_F \subseteq C(H_2)_F$, or
- (ii) $C(H_1)_T \subseteq C(H_2)_T$ and $C(H_1)_F \subset C(H_2)_F$

While the fundamental idea underlying P-Truthlikeness—namely that a theory’s truthlikeness depends on the content it shares or does not share with the whole truth—is both reasonable and *prima facie* plausible, this account has proven to be severely flawed. Tichý (1974) and Miller (1974) independently showed that P-Truthlikeness entails that no false theory can be more truthlike than any other false theory.⁵ This consequence is highly problematic as it shows that P-Truthlikeness cannot allow clarifying what it takes for a potential, yet incorrect, explanation of *e* to be more truthlike than another potential explanation of *e*.

Several philosophers argue that the Tichý-Miller paradox results primarily from the fact that, given Popper’s conception of a theory’s content, any of its logical consequences is deemed relevant to determine its truthlikeness.⁶ As Schurz and Weingartner (2010, p. 425) outline, if a true proposition *p* is part of $C(H)$ and a false proposition *q* is part of it too, then the false consequence $p \wedge q$ is also part of it and, therefore, *H*’s truthlikeness also depends on $p \wedge q$. Yet, $p \wedge q$ is merely a repetition of two of *H*’s content-parts. Likewise, if the true proposition *p* is part of $C(H)$ while the proposition $\neg r$ is not, then the true proposition $p \vee \neg r$ is also part of $C(H)$. However, $p \vee \neg r$ is only an arbitrary disjunctive weakening of *p* and can hardly be regarded as part of *H*’s semantic content. As Gemes (2007, p. 295) points out, if one regards $p \vee \neg r$ as part of *H*’s content when *p* is and $\neg r$ is not, one thereby commits to the view that part of *H*’s content is confirmed when new evidence confirming that $\neg r$ is acquired. Yet, such a view is quite implausible for *H* does not, in fact, say anything about $\neg r$. These considerations suggest that the very weakness of P-Truthlikeness is the notion of content which it depends on. According to this account, a theory’s truthlikeness depends on some of its logical consequences that are either only repetitions of some of its content-parts or arbitrary consequences of some of its content-parts.

Various proposals have been made to overcome the weakness of Popper’s account. Here, I will focus on Schurz and Weingartner’s (2010) proposal that has proven to have satisfying results compared to other solutions developed in the same spirit and which is paradigmatic of a range of solutions that are illuminating for the hypothesis under examination. They (2010, p. 426) propose to conceive of the content of a theory *H* as its class of relevant elementary consequences, written $E(H)$. The class of *H*’s relevant consequences is defined as follows, for any proposition *p* such that $p \in C(H)$:⁷

Relevant Consequence: *p* is a relevant consequence of *H* iff no propositional variable in *p* can be replaced on some of its occurrences by an arbitrary new variable *salva validitate* of the entailment $H \vdash p$.

This definition excludes from the class of consequences which H's truthlikeness depends on any arbitrary logical consequences of H, such as arbitrary disjunctive weakening of H's content-parts. As to the definition of H's elements, it excludes from the class of consequences any repetitions of H's logical consequences by decomposing H's consequences into their shortest conjunctive parts, where p_i is shorter than p_j just in case p_i is a well-formed formula that contains fewer logical operators than p_j . H's elementary consequences are defined as follows, for any proposition p such that $p \in C(H)$:

Element: p is an element of H iff p is not logically equivalent to a conjunction $p_1 \wedge \dots \wedge p_n$ ($n \geq 1$) in which each p_i is shorter than p .⁸

Schurz and Weingartner thus propose to represent the semantic content of a theory only by its relevant and elementary consequences. As they (2010, p. 427) demonstrate, $E(H)$ is logically equivalent to the theory H and is, therefore, able to represent its entire informational content. It is, however, important to note that one does not need to commit to the specific definitions offered by Schurz and Weingartner to accept the idea that a theory's content should be represented as its class of relevant and elementary consequences. While the results of Schurz and Weingartner's proposal are satisfying for propositional logic, Gemes (2007, pp. 304–305) shows that it has unwelcome consequences when applied to higher-order logics and offers an alternative proposal that avoids those consequences. Hence, while Schurz and Weingartner's approach is well-suited for my present purpose, other definitions of a theory's class of relevant elements might ultimately be deemed preferable.⁹

Based on the definition of a theory's class of relevant elements they offer, Schurz and Weingartner define what it takes for a theory to be more truthlike than another comparable theory as follows, where $E(H)_T$ is the class of H's true relevant elements and $E(H)_F$ its class of false relevant elements:

SW-Truthlikeness: $H_i >_T H_j$ iff:

- (i) $E(H_i)_T \vdash E(H_j)_T$ and $E(H_j)_T \not\vdash E(H_i)_T$ and $E(H_j)_F \vdash E(H_i)_F$, or
- (ii) $E(H_i)_T \vdash E(H_j)_T$ and $E(H_j)_F \not\vdash E(H_i)_F$ and $E(H_j)_F \vdash E(H_i)_F$.¹⁰

Before examining how this definition helps in clarifying the hypothesis under examination, let me come back to the notion of an explanation's content. In the previous chapter, I argued that J-Explanationism is not committed to the claim that for p to be part of an explanation H, p has to contribute to the explanation of e . In particular, I endorsed Popper's conception of an explanation's content and claimed that a proposition p can be part of H by virtue of the logical relation it bears to some of H's content-parts which themselves contribute to the explanation of e . Yet, the above considerations show that such a conception of an explanation's content is not a plausible

one. Arbitrary consequences of some of H's content-parts cannot themselves be regarded as being part of H's content although they are entailed by some of H's content-parts. As a result, I will henceforth restrict H's content to the class of H's relevant elementary consequences. As just noted, E(H) is sufficient to represent H's entire informational content. It should be clear, however, that conceiving of H's content in such a way does not commit one to the claim that p can be available to a subject as part of an explanation H only if p is a member of E(H). Indeed, the availability condition involved in J-Explanationism is not itself dependent on any specific conception of an explanation's content. Even if H's content is conceived of as H's class of relevant elements, a proposition which is not itself part of E(H) can be available to a subject as part of H by virtue of being available to her as a logical consequence of some members of E(H).

2.3. *Evidential Support and Expected Truthlikeness*

By defining the relative truthlikeness of a theory in terms of its true and false relevant and elementary consequences, SW-Truthlikeness shows how expectations for the relative truthlikeness of a theory amount to expectations for the truth or falsity of that theory's content. According to SW-Truthlikeness, a theory H_i is more truthlike than another theory H_j if and only if H_i 's class of true relevant elements (but not its class of false relevant elements) exceeds H_j 's class of true relevant elements, or H_j 's class of false relevant elements (but not its class of true relevant elements) exceeds H_i 's class of false relevant elements.¹¹ Consequently if H_i , given e , can be expected to be more truthlike than H_j , then either H_i 's class of true relevant elements (but not its class of false relevant elements) can be expected to exceed H_j 's class of true relevant elements, or H_j 's class of false relevant elements (but not its class of true relevant elements) can be expected to exceed H_i 's class of false relevant elements.

What still needs to be specified when it comes to the hypothesis under examination are the conditions under which the best explanation a subject has for e can be expected, given e , to be more truthlike than the other potential explanations she has *by virtue of containing the true proposition p* . This for any proposition p such that p is part of the best explanation a subject has for e in the sense specified at the end of the previous section. Indeed, it is only once those conditions have been specified that the truthlikeness hypothesis, understood in light of SW-Truthlikeness, can shed light on the general account of evidential support discussed in the previous chapter.

Under the assumption that p is part of a potential explanation H_i of e , H_i can be expected, given e , to be more truthlike than another potential explanation H_j of e because it contains the true proposition p only if two conditions are satisfied. First, p 's probability conditional on e has to be higher than $\neg p$'s probability conditional on e . The reason for this is that unless p is probable conditional on e , H_i 's class of true relevant elements cannot be

expected, given e , to contain the proposition p . Second, it has to be the case that H_1 's class of true relevant elements can be expected, given e , to exceed H_j 's class of true relevant elements by virtue of containing p .

In order for the second condition to be satisfied, p has to be distinctive of H_1 's semantic content relative to H_j . That is to say, p has to be a consequence of H_1 but not a consequence of H_j . To see why, suppose that H_1 entails the propositions p and q , while H_j only entails q . In such a case, H_1 's class of true relevant elements cannot exceed H_j 's class of true relevant elements by virtue of containing q as that proposition is also part of H_j 's semantic content. In contrast, H_1 's class of true relevant elements can exceed H_j 's class of true relevant elements by virtue of containing p as p is not part of H_j 's semantic content. This generalizes to any set of potential explanations available to a subject at a given time. Let $\{H_1 \dots H_n\}$ be that set and H_1 be the best explanation of that set such that $n \geq 2$ and $1 \leq i \leq n$. If H_1 can be expected to be more truthlike than any other member of that set, then for each explanation H_j member of that set such that $j \neq i$, H_1 can be expected to be more truthlike than H_j . But, for each potential explanation H_j , H_1 can only be expected to be more truthlike than H_j by virtue of containing the true proposition p if p is distinctive of H_1 's semantic content relative to H_j 's semantic content. Consequently, H_1 can only be expected to be more truthlike than the other explanation members of that set by virtue of containing, among other propositions, the true proposition p if p is a consequence of H_1 but is not a consequence of at least one of the other members of the set.

The best explanation a subject has for e can thus be expected, given e , to be more truthlike than the other potential explanations the subject has by virtue of containing the true proposition p only if p satisfies two conditions. First, p has to be probable conditional on e . Second, p has to be distinctive of that explanation's content relative to the other potential explanations available to the subject. Now, it is reasonable to think that these two conditions are conditions on the justification a subject can have for believing a proposition that is part of the best explanation she has for her evidence. This is because, plausibly, for any proposition p that is part of the best explanation a subject has for e , it is only when that explanation can be expected, given e , to be more truthlike than the other explanations available to her by virtue of containing the true proposition p that e supports believing p . Accordingly, the examination of the truthlikeness hypothesis suggests that a subject is justified in believing a proposition p that is part of the best explanation she has for e only if p is distinctive of that explanation's content and probable on the evidence possessed by the subject.

One important upshot of the first of the two conditions just put forward is that it shows precisely why probabilistic considerations are relevant to evidential support from an explanationist perspective. The second condition, for its part, follows quite naturally from that very perspective. If evidential support depends on explanatory considerations, it would be surprising that a subject with e was justified in believing that p in case p is neither a

distinctive content-part of the best explanation she has for e nor entailed by some distinctive content-parts of that explanation. For instance, suppose that a subject has two potential explanations for e — H_1 and H_2 —and that H_1 is a better potential explanation than H_2 . In addition, suppose that both of those explanations state that p and that p is not entailed by a proposition that is distinctive of H_1 's semantic content relative to H_2 's content. If evidential support depends on explanatory considerations, why would, in such a case, the subject be justified in believing that p ? As H_1 also states that p , the subject cannot be justified in believing that p because H_1 states that p and H_2 would explain her evidence better than H_1 .

One might worry, however, that the elucidation of those conditions, when taken as conditions on evidential support, is ultimately circular. For the truthlikeness hypothesis is itself a claim concerning what belief the evidence possessed by a subject supports. How can a claim which itself involves the notion of evidential support allow further explicating that very notion? In response to this particular worry, let me stress that the notion of evidential support has not been explicated, in the context of the present discussion, in terms of epistemic expectations concerning the relative truthlikeness of an explanation. Instead, the truthlikeness hypothesis which consists of a *prima facie* plausible claim pertaining to such expectations has been used to clarify the nature of the relation between explanatory considerations and expectations concerning the truth of an explanation's content. This, in turn, has allowed shedding better light on the conditions under which a subject's evidence supports believing a proposition that is part of the best explanation she has for her evidence.

Let me now come back to J-Explanationism. How should that account of evidential support be formulated in light of the discussion of the truthlikeness hypothesis? A first temptation is simply to require that p be a distinctive part of the best explanation the subject has for her evidence that is probable on that evidence. Given the conception of an explanation's content endorsed in the previous section, this would amount to ruling out many logical consequences of the best explanation a subject has for her evidence from the propositions that can be supported by that evidence. Yet, those logical consequences can also be supported by the evidence the subject possesses. If p is supported by a subject's evidence because the best explanation that subject has for her evidence can be expected to be more truthlike than its rivals by virtue of containing the true proposition p , then the evidence possessed by the subject supports any proposition q that logically follows from p . For, if q is a logical consequence of p , then, unless q is true, the best explanation the subject has for her evidence cannot be more truthlike than its rivals by virtue of containing the *true* proposition p . As a result, let me offer the following formulation of J-Explanationism:

J-Explanationism: Whenever p is available to S as part of an explanation for e , S , with e , is justified in believing that p iff either:

- (i) p is a distinctive part of the best explanation S has for e that is probable conditional on e , or
- (ii) p is a logical consequence of some content-parts of the best explanation S has for e which are themselves distinctive of that explanation and probable conditional on e .

Before examining this version of J-Explanationism in more detail, let me offer a final clarification regarding my discussion of the truthlikeness hypothesis. In case the best explanation a subject has for e can be expected to be more truthlike than the other potential explanations she has by virtue of containing the true proposition p , one may ask if forming the belief that p would allow that subject to better approximate the whole truth. While reasons have been provided by Niiniluoto (2004) and Cevolani (2013) to think that belief revision through abductive inferences allows one's belief system to better approximate the whole truth about a particular domain, it should be clear that forming beliefs that are supported by one's evidence does not always allow better approximating the whole truth about a particular domain. This is because the fact that some explanation H can be expected to be more truthlike than its rivals by virtue of containing the true proposition p entails neither that H is actually more truthlike than its rivals nor that p is true.

2.4. Explanationism

In addition to resulting from a plausible elucidation of the connection between explanatory considerations and the truth of an explanation's content, the version of J-Explanationism just put forward is able to account for the intuitions elicited by the Investigation case. Recall that this case represents a challenge because given the evidence Sally gathered during the first four steps of her investigative procedure, the proposition 'Jeremy committed the burglary' is part of the best explanation she has for her overall evidence. Yet, Sally does not appear to be justified in believing that Jeremy committed the burglary mid-way through her investigative procedure as it often happened that, after completing the last steps of the procedure, a new suspect emerged that better fitted the evidence.

According to the version of J-Explanationism offered in the previous section, p , to be supported by e , must either be a distinctive part of the best explanation a subject has for e that is probable on e or be available as a logical consequence of some of its content-parts which are themselves distinctive of that explanation and probable conditional on e . In the Investigation case, the proposition 'Jeremy committed the burglary' is part of the best explanation Sally has for her evidence. However, that explanation cannot be expected, prior to bringing Sally's investigative procedure to a conclusion, to be more truthlike than the other explanations available to Sally by virtue of containing the true proposition 'Jeremy committed the burglary.'

This is because this proposition's probability conditional on Sally's overall evidence is low. As the proposition 'Jeremy committed the burglary' is not available as a logical consequence of some other probable content-parts of the best explanation Sally has for her evidence, J-Explanationism delivers the result that Sally is not justified in holding the belief that Jeremy committed the burglary mid-way through her investigative procedure.

The reason why the version of J-Explanationism offered in the previous section is able to deliver the right result in the Investigation case is that, as already outlined, this account shows why and how probabilistic considerations are relevant to determine whether a subject is justified in believing a particular proposition given the evidence she has. One might worry, however, that as the notion of evidential probability plays a central role in the elucidation of the conditions under which some explanation can be expected to be more truthlike than its rival by virtue of containing the true proposition p , J-Explanationism suffers from the same weakness as the first hypothesis considered at the beginning of this chapter. More precisely, one might worry that given the version of J-Explanationism offered in the previous section, evidential support ends up simply being a matter of probabilistic considerations. Yet, this worry would be misplaced. While the best explanation a subject has for e can be expected to be more truthlike than its rivals by virtue of containing the true proposition p only if p is probable conditional on e , explanatory considerations remain central when it comes to determining whether a subject is justified in believing a particular proposition. For the fact that the best explanation a subject has for e can be expected to be more truthlike than its rivals ultimately depends on that explanation's merits.

In fact, J-Explanationism delivers the same results as the version of the account examined in the previous chapter in cases such as the Preface Paradox. In the Preface Paradox, Sam is justified in believing the conjunction of all the claims made in the body of his book although this conjunction is improbable on his overall evidence. As suggested by Cevolani and Schurz (2017, pp. 218–219), since the conjunction of the claims expressed in the body of Sam's book represents Sam's best attempt at approximating the truth regarding his domain of inquiry, the evidence gathered by Sam supports believing that conjunction. The account offered in his book can be expected, given the evidence Sam gathered, to be more truthlike than its rivals by virtue of containing the claims expressed in the body of the book and, as a result, Sam is justified in believing what logically follows from those claims even if his evidence also supports believing that the chance that the body of the book contains at least one mistaken claim is high.

The results delivered by J-Explanationism in the Preface Paradox illustrate the fact that according to this account any logical consequence of the best explanation a subject has for her evidence can, in principle, be supported by that evidence. As a matter of fact, J-Explanationism is able to deliver the correct results in standard Gettier cases which typically involve a valid inference

to an arbitrary consequence of some content-parts of the best explanation a subject has for her evidence. Consider Gettier's (1963, pp. 122–123) example in which Smith has evidence for the proposition 'Jones owns a Ford' and infer from this proposition that either Jones owns a Ford or Brown is in Barcelona. His evidence includes that, in his memory, Jones always owned a Ford and that Jones just offered him a ride while driving a Ford. Now, even if the proposition 'either Jones owns a Ford or Brown is in Barcelona' is merely an arbitrary disjunctive weakening of a content-part of the best explanation Smith has for his evidence, J-Explanationism delivers the result that Smith is justified in believing that Jones owns a Ford or Brown is in Barcelona.¹² The proposition 'Jones owns a Ford or Brown is in Barcelona' is available to him as a logical consequence of a proposition—*i.e.* 'Jones owns a Ford'—that is part of the best explanation he has for his evidence. I will come back, later in this book, to Gettier cases and, in particular, to the relation between doxastic justification and knowledge. For now, let me simply stress that the ability of J-Explanationism to deliver intuitive results in standard Gettier cases makes it clear that, according to this account, a subject can be justified in believing the true proposition p without knowing that p .

2.5. Concluding Remarks

The Investigation case discussed in Chapter 1 raised an important challenge for the version of J-Explanationism that was put forward in that chapter, for, according to it, probabilistic considerations did not seem to have any justificatory import. To address this problem, I proposed to bracket the approach of evidential support discussed in Chapter 1 to examine a *prima facie* plausible claim concerning the connection between explanatory considerations and the truth of an explanation's content. According to that claim—the truthlikeness hypothesis—the best explanation a subject has for e can be expected, given e , to be more truthlike than the other explanations that subject has for e . Given a plausible conception of an explanation's content and of what it takes for some explanation to be more truthlike than another comparable explanation, the truthlikeness hypothesis allowed clarifying the connection between explanatory considerations and the truth of an explanation's content as well as the conditions under which, from an explanationist perspective, a subject's evidence supports believing a proposition. This, in turn, allowed formulating an updated version of J-Explanationism which shows how probabilistic considerations bear on evidential support and, as a result, is able to deliver the right result in the problematic cases examined in Chapter 1.

Notes

- 1 This hypothesis has been defended by Lipton (2003, 2004) with respect to abductive inferences. According to him, the potential ability of an explanation

- to provide an understanding of why e is a good guide to estimate its probability conditional on e .
- 2 Kuipers (2004) and Niiniluoto (2004, 2005) both worked on the relation between theories' expected degree of truthlikeness and abductive inferences in order to formulate rules of abductive inference that conform to our intuitions regarding the ways subjects should revise their beliefs. The rule formulated by Kuipers prescribes one to infer that the best available explanation is more truthlike than any other available potential explanation and relies, therefore, on the same plausible connection I just stressed between explanatory considerations and the relative truthlikeness of potential explanations.
 - 3 Note that although I proposed to bracket the account of evidential support discussed in the previous chapter, the truthlikeness hypothesis is plausible by that account's own lights. For the claim that H is more truthlike than its rivals best explains why H can provide a better answer to the question "why is e the case?" than its rivals.
 - 4 This chapter focuses exclusively on definitions of comparative truthlikeness and not on quantitative measures of truthlikeness that have been offered, *inter alia*, by Niiniluoto (2011) and Cevolani and Schurz (2017). The reason for this focus is that, as just noted in the previous section, the truthlikeness hypothesis relies on a comparative notion of truthlikeness.
 - 5 This result is due to the fact that if the falsity-content of a theory H is not empty, H 's truth-content cannot be increased without thereby increasing H 's falsity-content. By adding a true proposition p to $C(H)$ one, thereby, adds a false proposition $p \wedge q$ to $C(H)$ where q is a false consequence of H already contained in $C(H)$.
 - 6 See Mott (1978), Schurz and Weingartner (1987, 2010), Burger and Heidema (1994), Gemes (2007), and Cevolani, Festa, and Kuipers (2013) for later definitions of a theory's truthlikeness based on this observation.
 - 7 Schurz and Weingartner (2010, pp. 425–426) propose as conventions that all members of $E(H)$ be written in their negation normal form—*i.e.*, by writing them only with the logical operators " \wedge , \vee , \neg ", by putting all negations in front of propositional variables, and by eliminating all double negations and associative brackets—and that all disjunctions of literals (propositional variables or their negations) in $E(H)$ be written as clauses—*i.e.*, as disjunctions of literals in distinct and alphabetically ordered propositional variables.
 - 8 Note that, according to this definition, a proposition p such that $p \in C(H)$ can be an elementary consequence of H without being a relevant consequence of H . Likewise, p can be a relevant consequence of H without being an elementary consequence of H .
 - 9 See also Cevolani and Festa (2018) for a similar approach of truthlikeness.
 - 10 SW-Truthlikeness is formulated in terms of entailment of classes of relevant elements instead of inclusion because classes of relevant elements are not themselves closed under logical consequence.
 - 11 $E(H_i)_T$ exceeds $E(H_j)_T$ whenever some members of $E(H_i)_T$ are not deducible from $E(H_j)_T$ but all members of $E(H_j)_T$ are deducible from $E(H_i)_T$. In other words, $E(H_i)_T$ exceeds $E(H_j)_T$ just in case all true relevant elements of H_j are part of H_i 's semantic content but some of H_i 's true relevant elements are not part of H_j 's semantic content. The same holds for $E(H_i)_F$ and $E(H_j)_F$.
 - 12 The other case offered by Gettier involves the formation of a lucky true belief concerning an existential claim that is inferred from a non-quantified proposition. While Schurz and Weingartner's (2010) definition of a theory's relevant consequences class does not deal with quantified propositions, Gemes's (2007) definition, for instance, which is developed for higher-order logics, typically excludes such consequences from the class of H 's relevant consequences.

References

- Burger, I. C. & Heidema, J. (1994). Comparing Theories by their Positive and Negative Contents. *British Journal for the Philosophy of Science*, 45, 605–630.
- Cevolani, G. (2013). Truth Approximation via Abductive Belief Change. *Logic Journal of IGPL*, 21, 999–1016.
- Cevolani, G. & Festa, R. (2018). A Partial Consequence Account of Truthlikeness. *Synthese*, 197, 1627–1646.
- Cevolani, G., Festa, R. & Kuipers, T. (2013). Verisimilitude and Belief Change for Nomic Conjunctive Theories. *Synthese*, 190, 3307–3324.
- Cevolani, G. & Schurz, G. (2017). Probability, Approximate Truth, and Truthlikeness: More Ways Out of the Preface Paradox. *Australasian Journal of Philosophy*, 95, 209–225.
- Foley, R. (1992). The Epistemology of Belief and the Epistemology of Degrees of Belief. *American Philosophical Quarterly*, 29, 111–121.
- Foley, R. (2009). Beliefs, Degrees of Belief, and the Lockean Thesis. In H. Franz & C. Schmidt-Petri (Eds.). *Degrees of Belief*. Dordrecht: Springer, 37–47.
- Gemes, K. (2007). Verisimilitude and Content. *Synthese*, 154, 293–306.
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis*, 23, 121–123.
- Kuipers, T. (2004). Inference to the Best Theory, Rather Than Inference to the Best Explanation – Kinds of Abduction and Induction. In F. Stadler (Ed.). *Induction and Deduction in the Sciences*. Dordrecht: Springer Science, 25–51.
- Lipton, P. (2003). Is Explanation a Guide to Inference? A Reply to Wesley Salmon. In G. Hon & S. Rakover (Eds.). *Explanation: Theoretical Approaches and Applications*. Dordrecht: Springer, 93–120.
- Lipton, P. (2004). *Inference to the Best Explanation*. London: Routledge.
- Miller, D. (1974). Popper’s Qualitative Theory of Verisimilitude. *The British Journal for the Philosophy of Science*, 25, 166–177.
- Mott, P. (1978). Verisimilitude by Means of Short Theorems. *Synthese*, 38, 247–273.
- Niiniluoto, I. (2004). Truth-Seeking by Abduction. *Vienna Circle Institute Yearbook*, 11, 57–82.
- Niiniluoto, I. (2005). Abduction and Truthlikeness. In R. Festa, A. Aliseda & J. Peijnenburg (Eds.). *Confirmation, Empirical Progress, and Truth Approximation: Essays in Debate with Theo Kuipers*. Amsterdam: Rodopi, 255–275.
- Niiniluoto, I. (2011). The Development of the Hintikka Program. In D. M. Gabbay, J. Woods & S. Hartmann (Eds.). *Handbook of the History of Logic: Inductive Logic*. Amsterdam: Elsevier, 311–356.
- Popper, K. R. (1963). *Conjectures and Refutations*. London: Routledge.
- Schurz, G. & Weingartner, P. (1987). Verisimilitude Defined by Relevant Consequence-elements. In T. A. Kuipers (Ed.). *What is Closer-to-the-truth?* Amsterdam: Rodopi, 47–78.
- Schurz, G. & Weingartner, P. (2010). Zwart and Franssen’s Impossibility Theorem Holds for Possible-world-accounts but not for Consequence-accounts to Verisimilitude. *Synthese*, 172, 415–436.
- Tichý, P. (1974). On Popper’s Definitions of Verisimilitude. *The British Journal for the Philosophy of Science*, 25, 155–160.

3 On Understanding

3.1. *The Grasping Requirement*

If the evidence a subject possesses can be conceived of as what is in need of an explanation for that subject, it is reasonable to think that doxastic justification is closely connected to understanding and, in particular, to explanatory understanding.¹ Explanations are after all sets of propositions which can, by virtue of the structure they instantiate, provide one with understanding and if J-Explanationism is correct, it is plausible that the propositions a subject is justified in believing given her overall evidence are the ones that can put her in a position to understand at least some aspects of that evidence.²

More precisely, let e_H stand for the portion of a subject's overall evidence that could be explained by H.^{3,4} If a subject, given e , is justified in believing the content-parts of H, then one might think that she is thereby in a position to understand why e_H . This is because, if the subject is justified in believing H's content-parts, then H consists of the best explanation she has of why e_H is the case. Given e , H represents her best attempt at accounting for why e_H . Yet, as outlined in Chapter 1, the conditions under which a proposition p is available to a subject as part of an explanation H of her evidence should not be construed as being overly demanding. This means, essentially, that the fact that a subject is justified in believing the content-parts of H does not guarantee that she is in a position to account for e_H by means of the propositions she is justified in believing. Even if H is the best explanation the subject has for her evidence, she might fail to grasp H in such a way as to be able to account for e_H by means of that explanation.

Several philosophers acknowledge that understanding involves a grasping requirement.⁵ According to Kvanvig (2003, p. 192), for instance, understanding requires "the grasping of explanatory and other coherence-making relationships in a large and comprehensive body of information." Elgin, considering the understanding one can have of the Athenian victory over Persia at Marathon, states for her part that:

Understanding the Athenian victory involved more than knowing the various truths that belong to a suitable tethered comprehensive, coherent account of the matter. The understander must also *grasp* how the various truths are related to each other and to other elements of the account.

(Elgin, 2017, p. 46)

Now, p can be available to a subject as part of H in the sense specified in Chapter 1 without that subject grasping H 's content in the required way. To be sure, consider the following example due to Pritchard:

Faulty Wiring: Suppose that I understand why my house burned down, know why it burned down, and also know that it burned down because of faulty wiring. Imagine further that my young son asks me why his house burned down and I tell him. He has no conception of how faulty wiring might cause a fire, so we could hardly imagine that merely knowing this much suffices to afford him understanding of why his house burned down. Nevertheless, he surely does know that his house burned down because of faulty wiring, and thus also knows why his house burned down.

(Pritchard, 2010, p. 81)

A correct explanation of why the house burned down is made available to the son and the son, assuming he has the conceptual resources to retrieve the meaning of what is said to him, is disposed to treat the content of that explanation as part of an answer to the question “why did the house burn down?” Yet, as he lacks a conception of how faulty wiring might cause a fire, he does not appear to be able to grasp that explanation in such a way as to understand why the house burned down by means of it. Presumably, he can come to believe, with justification, that what is said to him explains why the house burned down and he might offer that explanation as an answer to the question “why did the house burn down?” But he is unable to see how the explanation, given its content, can account for the fact that the house burned down.⁶

3.2. *Grasping and Knowing*

Thus, for a subject to be in a position to understand why e_H , it is not sufficient that H be available to her as an explanation of e_H . The subject must grasp how H can, given its content, account for e_H . But what does it take for a subject to have such a grasp?

If explanations can be conceived of as sets of propositions which instantiate a certain structure, a natural thought is that the grasping requirement of understanding pertains to the structure instantiated by a given explanation.

According to Strevens, for instance, understanding a given phenomenon by means of an explanation requires grasping:

That the states of affairs represented by the propositions in fact obtain, and second, that the propositions instantiate the prescribed structure—for example, that they form a deductive argument for the explanandum (for Hempel) or that they stand in the right kinds of statistical relationship to the explanandum and to each other (for Salmon).
(Strevens, 2013, p. 511)

What is required of the son who learns that the house burned down because of faulty wiring is, according to this proposal, that he grasps that the states of affairs represented by the explanation provided to him obtain and that the propositions that are part of that explanation instantiate a certain structure. In Strevens's view, by grasping this, the son understands that the answer provided to him explains why the house burned down.

Note that this conception of the grasping requirement of understanding involves two distinct elements. First, an element related to the obtaining of the states of affairs represented by the explanation. Second, an element related to the structure of the explanation labeled the *narrow* or *psychological* component of grasping by Strevens (2013, p. 511). He regards these two elements as being separable in that he recognizes two possible notions of grasping; the first being characterized in terms of both elements and the second being characterized only in terms of the narrow component of grasping. It is, however, not clear that the first notion of grasping considered by Strevens is able to capture adequately the grasping requirement of understanding for there are reasons to consider that the first element of the proposal goes beyond this requirement. In the Faulty Wiring case, the son is plausibly justified in believing that the states of affairs represented by the explanation provided to him obtain and in that “grasps” or rather is in a position to know that they obtain. Yet, the specific kind of grasp involved in understanding is distinct from this requirement. A subject can grasp how an incorrect explanation of e_H that she has no reason to endorse can account for e_H . She can see how it can account for e_H independently of the evidence she has for believing that the states of affairs represented by the explanation obtain.

What Strevens labels the narrow component of grasping is, therefore, sufficient to capture the type of grasp required for the son to understand why the house burned down by means of the explanation provided to him. He needs to grasp how that explanation can account for the fact that the house burned down and having such a grasp plausibly amounts to grasping the structure instantiated by the explanation provided to him. Consequently, let me examine in more detail what it takes for a particular subject to grasp that structure. According to Newman (2012, 2013, 2017), a subject can come to understand an explanation as opposed to merely knowing it by performing certain inferences on each part of the explanation. Those inferences can be,

inter alia, logical, causal, goal-directed, or probabilistic and typically allow the subject to make adequate conceptual connections between the various elements of the explanatory story.⁷ In the case of a causal explanation H of e_H , those inferences are inferences to the reasons that explain why the causes cited in H are causally responsible for e_H and more generally:

For an agent A to understand an explanation of some phenomenon P, A must generate inferential knowledge of the reasons that are causally, logically, or probabilistically responsible for c being responsible for P.
(Newman, 2013, p. 178)

In Newman's view, it is thus by inferring specific pieces of explanatory knowledge that a subject can come to make adequate connections between the various elements of an explanation and thereby grasp its structure. It is by inferring the reasons that explain why faulty wiring caused the house to burn down that the son can grasp the explanation that is provided to him in the required way. This is because, to grasp how that explanation can account for the fact that the house burned down, the son must infer the reasons that are causally responsible for the faulty wiring causing the fire.

Newman's inferential model is illuminating when it comes to the elucidation of the narrow component of grasping identified by Strevens as it is plausible that a subject comes to make adequate connections between the elements of an explanation by performing certain inferences on each part of the explanation. However, I doubt that the kind of inferential knowledge identified by Newman is the relevant one when it comes to the grasping requirement of understanding.

Newman's proposal first raises a question given the very distinction it is supposed to elucidate: the distinction between understanding an explanation and merely knowing it. According to the proposal, it is by generating inferentially specific pieces of explanatory knowledge that a subject comes to understand an explanation H of e_H . The knowledge generated by the subject is explanatory because it concerns the reasons that *explain* why c —some element cited in H—is responsible for e_H . Accordingly, it is by inferring further explanations that a subject comes to understand or grasp H. But one can legitimately ask whether merely knowing the explanations that are inferred is sufficient for the subject to understand H. After all, Newman is working under the assumption that knowing an explanation importantly differs from understanding it. Part of the difference, as he (2017, p. 575) rightly notes, is that understanding an explanation involves knowing how to use it. It is, therefore, not clear that merely knowing the explanation of why c is responsible for e_H is sufficient to understand H. Should the subject who infers that explanation not also be in a position to *use* it to understand H?

This particular question is, however, not a reason, on its own, to doubt that Newman's inferential model adequately captures the grasping requirement of

explanatory understanding. The reason is, instead, that generating the type of inferential knowledge identified by Newman is unnecessary to acquire a grasp of how some explanation can account for e_H . As Pritchard (2014, p. 331) notes, to understand why the house burned down by means of the explanation provided to him, the son needs a grip on *how* the cause generated the effect. Of course, the terms “how” and “why” can often be substituted in questions. Asking “why does lighting occur?” does not appear to be substantially different from asking “how does lighting occur?” Nonetheless, providing an answer to a particular how-question does not always amount to providing an answer to the corresponding why-question. For instance, inferring that faulty wiring produced an electrical spark that inflamed the dry materials surrounding it would correctly answer the question as to how faulty wiring generated fire without answering the question as to why it did. What explains why faulty wiring caused the fire are properties of faulty wiring causally responsible for it producing an electrical spark that inflamed the dry materials surrounding it. The answer provided to the how-question fails to constitute an answer to the corresponding why-question because the former only describes the mechanism at work in the production of the fire by faulty wiring without identifying the properties of faulty wiring that are causally responsible for it being responsible for the fire.⁸

According to Newman’s model, to understand why the house burned down, the son has to ask *why* faulty wiring generated fire and infer the reasons that are causally responsible for faulty wiring being responsible for the fire. Yet, this is not required in order for the son to make the adequate connections between the elements of the explanation provided to him. To make those connections, it is sufficient for the son to ask *how* faulty wiring generated a fire. To be sure, let me consider a slightly different case. Suppose that a student who aims at understanding why mammals need oxygen is told by her biology teacher that the cells of mammals use oxygen to perform their functions. To understand why mammals need oxygen by means of this explanation, the student has to answer the question as to how the cells of mammals use oxygen to perform their functions. Now, to do so it is sufficient to infer that, for instance, oxygen enters into the production of adenosine triphosphate (ATP) molecules that enable the cells to perform their functions. Such an inference would not constitute an inference to the reason why oxygen is used by the cells of mammals to perform their functions but instead to how it is used. To know why oxygen is used by the cells of mammals to perform their functions, the student would need to draw on an extensive knowledge of the citric acid cycle to identify the properties of oxygen that are responsible for it being used by the cells of mammals in the way it is used. But answering the question as to how the cells of mammals use oxygen to perform their functions is all that is required for the student to make the adequate connections and grasp how the explanation provided to her can account for the fact that mammals need oxygen.

3.3. Knowledge of Dependencies

The grasping requirement of explanatory understanding is not essentially related to the type of inferential explanatory knowledge identified by Newman.⁹ To grasp how H can account for e_H , the inferences performed by a subject do not need to “reveal the reasons that the causal and logical properties in the explanation are the causal or logical properties that they are” (Newman, 2012, p. 15). But this, of course, raises the question as to what, precisely, those inferences need to reveal.

To address this question, let me outline that there exists an intuitive connection between one’s ability to apprehend various elements as depending on each other and one’s ability to understand why something is the case. Greco expresses this idea by characterizing understanding as a way of seeing how different things fit together, which, in his view, amounts to having a systematic knowledge of dependence relations:¹⁰

Understanding consists in a systematic knowledge of dependence relations, where dependence relations can be of various sorts, including ‘real’ relations between parts of the world, conceptual and logical relations between parts of a theory, and semantic relations between theory and world.

(Greco, 2014, p. 293)

While I agree with Greco regarding the connection between understanding and a subject’s apprehension of various elements as depending on each other, the considerations put forward in the previous section suggest that the inferences performed by a subject on each part of an explanation should not merely reveal that e_H depends on c according to some explanation H. Grasping H requires acquiring a knowledge of the *way* in which e_H depends on c according to the explanation and this, I submit, is precisely what the inferences performed on each part of the explanation can reveal to a subject.¹¹

Here, by “the way in which e_H depends on the elements cited in an explanation” I do not mean the type of dependence relation that e_H bears to those elements. Performing inferences that merely reveal that, according to H, e_H depends causally as opposed to logically on c , would only allow a subject to grasp what kind of explanation H is. What I mean is more specific. The relevant inferences performed by the student on each part of the explanation provided to her do not merely reveal that the performance of certain functions by the cells of mammals depends causally on their intake of oxygen. Instead, they reveal that oxygen entering into the production of ATP molecules is the way in which the performance of their functions by the cells of mammals depends on their intake of oxygen. As outlined in the previous section, however, the way in which the two elements of the explanation depend on each other does not need to be so specific as to make explicit

the reasons why they depend on each other in such a way. Grasping only requires generating inferential knowledge of the *ways* in which the elements of the explanatory story depend on each other and it is perfectly possible for someone to know that w is the way in which e_H depends on c without knowing why e_H depends on c in such a way.

Grimm's (2014) account of understanding relies on similar considerations. In his view, explanatory understanding consists in a knowledge of causes or, more generally, a knowledge of dependence relations.¹² However, according to Grimm this knowledge of dependence relations should not be conceived of on the model of propositional knowledge as this conception is precisely what generates the type of problems identified by philosophers such as Pritchard (2010, 2014) for the view that understanding amounts to knowing certain propositions.

Taking as an example the knowledge one can have of necessities such as '2+2=4,' Grimm (2014, p. 334) claims that one's *a priori* knowledge of such necessities typically results from one's apprehension of the necessity itself. One comes to know that 2+2=4 by seeing that this addition could not possibly have a different result, this seeing corresponding to the rational insight invoked by philosophers such as Bonjour (2005) to account for *a priori* justification. As a result, the primary object of such knowledge is not a proposition but modal reality itself:

The basic idea here is therefore not that propositions have no role to play in *a priori* knowledge, but rather that they play a secondary or derivative role. If Bonjour is right – and I think he is – the primary object of *a priori* knowledge is the modal reality itself that is grasped by the mind, and it is on the basis of this grasp that we then (typically) go on to assent to the proposition that describes or depicts these relationships.
(Grimm, 2014, p. 335)

According to Grimm, the kind of knowledge which explanatory understanding amounts to is similar to *a priori* knowledge as it involves a direct apprehension or grasp of modal reality. Its primary objects are, therefore, not propositions depicting dependence relations but the way things “stand in the modal space” (Grimm, 2014, p. 334). It is by grasping modal reality itself that a subject gains explanatory understanding and goes on to assent to propositions that depict dependence relations between the elements cited in an explanation.

Explanatory understanding, on this view, is thus a type of knowledge that is gained on the basis of a direct grasp of the way things stand in the modal space. As this knowledge consists of a knowledge of dependence relations, what needs to be grasped by a subject is the way in which various elements depend on each other. Indeed, Grimm does not simply claim that understanding why e_H involves knowing that e_H depends on c . Instead, understanding why e_H requires grasping how e_H and c stand in the modal space

and this amounts to grasping the way in which e_H and c depend on each other. Accordingly, Grimm's proposal also identifies an essential connection between the grasping requirement of understanding and one's ability to identify the ways in which the elements cited in an explanation depend on each other. Nevertheless, Grimm conceives of this ability as a direct insight into modal reality and it is far from clear that such a radical conception of the grasping requirement of understanding is needed.¹³

First, as pointed out by Khalifa (2017), there is some unclarity regarding the precise nature of Grimm's proposal. While Grimm states that explanatory understanding should not be conceived of on the model of propositional knowledge, he also suggests (2014, p. 334) that explanatory understanding, as *a priori* knowledge of necessities, is a type of *de re* knowledge that is gained on the basis of a grasp of modal reality. Yet, as Khalifa notes, the distinction that can be drawn between *de re* and *de dicto* knowledge is orthogonal to the distinction that can be drawn between propositional and non-propositional knowledge. It is, therefore, unclear why the conception of the grasping requirement of understanding put forward by Grimm should speak against a conception of understanding on the model of propositional knowledge even if this requirement is itself taken to be non-propositional.

Second, the reason why Grimm (2006, 2014) invokes a direct grasp of modal reality as a basis of the knowledge which understanding consists in is that, in his view, abilities to answer counterfactual questions about the understood phenomenon, and to anticipate changes in it, are distinctive of understanding. As he takes this direct grasp of modal reality to involve such abilities, the fact that understanding is a type of knowledge based on such a grasp explains why these abilities are distinctive of understanding. Yet, a direct grasp of modal reality is not required to account for these distinctive aspects of understanding. By performing the relevant inferences on each part of an explanation H , a subject can come to know the ways in which the elements cited in H depend on each other according to H . In addition, such knowledge is precisely what can allow a subject, upon her endorsement of H as the explanation of e_H , to answer counterfactual questions about e_H and to anticipate changes in it. The student's knowledge that oxygen entering into the production of ATP molecules is the way in which the performance of their functions by the cells of mammals depends on their intake of oxygen is what can allow her, upon endorsing the explanation provided to her, to answer questions such as "what if oxygen was lacking in the cells of a particular mammal?" It is because the student identified how the performance of their functions by the cells of mammals depends on their intake of oxygen that she is able to state that ATP molecules could not be produced in the cells of that particular mammal and that, as a result, its cells could not perform their functions. But inferring the ways in which the elements cited in an explanation depend on each other does not consist in having a *direct* grasp of modal reality although what is inferred by the subject pertains to the way things stand in the modal space according to the

explanation. By performing such inferences, a subject acquires knowledge that w is the way in which e_H depends on c according to an explanation H and such knowledge is precisely what puts the subject in a position to *use* H , upon endorsing it as the explanation of e_H , to answer counterfactual questions about e_H and anticipate changes in e_H .

3.4. *Grasping Explanations*

Grasping how an explanation H can account for e_H requires generating a specific type of inferential knowledge. It requires generating inferential knowledge of the ways in which the elements of the explanatory story depend on each other according to H . Consequently, the *act* of grasping an explanation is fundamentally inferential and should not be conceived of as a direct insight into the way things stand in the modal space. It provides an insight into modal reality by delivering knowledge that is propositional in nature. By grasping an explanation H , a subject comes to know that w is the way in which things depend on each other according to H and such knowledge is what puts the subject in a position to use H , upon endorsing it, to assess counterfactual scenarios pertaining to the explanandum. Indeed, the knowledge one gains by performing the relevant inferences on each part of an explanation has an important practical dimension. By gaining inferential knowledge of the ways in which the elements of the explanatory story depend on each other, one acquires a *knowledge of how* to account for the explanandum by means of the explanation and, as Grimm suggests, such knowledge involves knowing how to assess counterfactual scenarios pertaining to the explanandum by means of that explanation.¹⁴ By learning that oxygen entering into the production of ATP molecules is the way in which the performance of their functions by the cells of mammals depends on their intake of oxygen, the student essentially gains a knowledge of how to account for the fact that mammals need oxygen. She comes to know the way in which that explanation can be used to account for the fact that mammals need oxygen. Likewise, by learning that the production of an electrical spark that inflamed the dry material surrounding it is the way faulty wiring generated a fire, the son gains a knowledge of how to account for the fact that the house burned down. Thus, while the act of grasping an explanation is adequately conceived of as an inferential process, the state of grasping an explanation can be conceived of as a particular kind of know-how. As suggested by Elgin (2017, p. 33) who notes that “to grasp a proposition or an account is at least in part to know-how to wield it to further one’s epistemic ends,” having a grasp of an explanation H of e_H consists in knowing how H can be used to account for e_H .

Let me stress that this elucidation of the grasping requirement of explanatory understanding relies on an epistemological claim concerning the way a subject can acquire a particular kind of know-how rather than on the following ontological claim defended by Kim:

Explanations track dependence relations. The relation that ‘grounds’ the relation between an explanans, G, and its explanatory conclusion, E, is that of dependence; namely, G is an explanans of E just in case e, the event being explained, depends on g, the event invoked as explaining it. (Kim, 1994, p. 68)

According to Kim, relations of dependence ground explanatory relations in the sense that some phenomenon is explained by a correct explanation H by virtue of bearing some kind of dependence relations to the elements cited in H. The proposed elucidation of the grasping requirement of understanding does not, however, depend on a claim about the nature of explanations. Instead, it depends on the epistemological claim that it is by inferring how the elements cited in H depend on each other according to H that a subject can come to know how to account for a given phenomenon by means of H. While the ontological claim defended by Kim, if true, would add support to and explain the epistemological claim which I rely on, these two claims differ importantly. The epistemological claim could be true although Kim is incorrect concerning the nature of explanations. Moreover, the former claim has a broader domain of application than the latter. If, for instance, H is only a potential explanation of e_H , a subject can come to know how to use H to account for e_H by inferring how the elements cited in H depend on each other according to H. However, the relation between H and e_H cannot be grounded in relations of dependence between the elements of H and e_H as H does not actually explain why e_H .

Kim’s discussion of the nature of explanations is nevertheless illuminating when it comes to the contribution of the grasping requirement to a central feature of understanding. As he notes, one important upshot of his analysis of explanations is that it accounts for one of their key aspects highlighted by philosophers such as Friedman (1974) and Kitcher (1981):

The ontological contribution of dependence relations lies exactly in this fact: they reduce the number of independent events, states, facts, and properties we need to recognize. And that is precisely the unifying and simplifying power of dependence relations. Unity and structure go hand in hand; dependence enhances unity by generating structure. (Kim, 1994, p. 68)

Similar considerations speak for the conclusion that because of its grasping requirement, understanding involves acquiring a more unified and structured picture of reality. According to what I argued, knowing how to account for e_H by means of H requires gaining inferential knowledge of the ways in which the elements cited in H depend on each other according to H. If this is correct, then knowing how to account for e_H by means of H requires discovering ways in which the number of independent events, states, facts and properties can be reduced by using H. Accordingly, understanding why e_H

by means of H involves acquiring a more unified and structured picture of reality. By understanding why e_H by means of H, a subject reduces the number of independent events, states, facts, and properties. But knowing how to use an explanation to account for a phenomenon does not merely involve knowing that, according to that explanation, the number of independent events, states, facts, and properties can be reduced. It involves having a conception of the ways in which the elements of the explanatory story depend on each other. For this reason, understanding does not only involve acquiring a more unified and structured picture of reality. It also involves acquiring a picture of reality that is such that changes in that reality can be anticipated. Because of its grasping requirement, the unifying force of understanding thus coincides with the abilities that are distinctive of that cognitive standing such as the ability to answer what Woodward (2003) labels “what-if-things-were-different?” questions.

3.5. *Grasping and Justification*

Understanding why e_H by means of an explanation H requires that the subject knows how to account for e_H by means of H. It requires that she knows how e_H depends on the elements cited in H according to that explanation. Consequently, as noted in the first section of this chapter, the fact that a subject is justified in believing the content-parts of H cannot guarantee that she is in a position to understand why e_H by means of H. Even if H is available to the subject and represents, given e , her best attempt at accounting for e_H , she can lack the know-how required to understand why e_H by means of H. Does it follow that *in addition* to being justified in believing the content-parts of H a subject must know how to account for e_H by means of H in order for her to be in a position to understand why e_H ?

According to Hills (2009, 2016) and Dellsén (2017, 2018, 2019), understanding, in fact, does not require having justification for believing propositions relevant to the explanation of the target phenomenon. Consider the following case discussed by Hills:

Napoleon: Suppose that you read in your book that Napoleon was tactically astute, and so on, and on the basis conclude that he was a great leader. But now your history teacher, whom you regard as extremely trustworthy, tells you that Napoleon was not a great leader. Your teacher is not basing this judgement on other information or on a different interpretation of what it takes to be a great general: he simply irrationally dislikes Napoleon. You have no idea about any of this, but even so, you ignore your teacher and continue to maintain your conclusion. Just as in the previous example, you have the abilities required for understanding, your beliefs are correct and in short, you understand why Napoleon was great.

(Hills, 2016, p. 672)

In Hills's view, the fact that the teacher's testimony defeats the justification the subject has for believing that Napoleon was a great leader does not prevent her from understanding why Napoleon was tactically astute and so on. The subject grasps the explanation in the way required to understand what she read in her book and the piece of defeating evidence she acquires, while defeating her justification for believing that Napoleon was a great leader, has no impact on the understanding she has of the facts about Napoleon reported in her book.

Contrary to what Hills suggests, however, I believe that a subject being in a position to understand why something is the case depends on the justification she has for endorsing the content of an explanation. To see why let me first examine how a subject's grasp of an explanation is reflected in the understanding she gains by means of that explanation. According to Kelp (2015, p. 3810), a subject understands a phenomenon P "when the basing relations that obtain between the agent's beliefs about P reflect the agent's knowledge about the explanatory and support relations that obtain between the members of the full account of P." In other words, when a subject understands why e_H by means of H, the way her beliefs concerning the propositions that are available to her as part of H are based reflects the subject's grasp of H. This is particularly plausible once the grasping requirement of understanding is conceived of as I argued it should in the previous sections. If to grasp H in the required way is to know how to account for e_H by means of H, then, plausibly, when a subject understands, she believes the propositions that are available to her as part of H because of the way H could account for e_H . That is, when a subject understands why e_H by means of H, her knowledge of how to account for e_H by means of H plays an essential role in her commitment toward H. In the Faulty Wiring case, for instance, the son endorses the content of the explanation that is provided to him. Yet, the way his beliefs are based doesn't reflect a grasp of that explanation. As he does not know how the fire depends on the elements cited in the explanation provided to him, the son has no choice but to rely solely on someone's testimony when it comes to endorsing the content of that explanation. He cannot be regarded as believing that content because of the way the explanation could account for the fact that the house burned down.

If I am correct, understanding why e_H by means of H thus involves endorsing H's content because of the way H could account for e_H . Yet, even if the subject's knowledge of how to account for e_H by means of H plays an essential role in the subject's commitment toward H, that commitment amounts to no more than a guess if, given the subject's overall evidence, she cannot expect H to be the right way to account for e_H . Indeed, if a subject cannot expect H to be the right way to account for e_H , then, from the perspective of her evidential situation, it is a matter of luck whether the way in which H could account for e_H is actually the right way to account for e_H . As a result, even if she endorses H's content because of the way H could account for e_H , her endorsement of H only amounts to taking a (possibly

lucky) guess regarding the right way to account for e_H . But understanding cannot plausibly result from taking such a guess. To be in a position to understand why e_H by means of H, the subject's evidential situation must be such that, given her overall evidence, H can be expected to be the right way to account for e_H .

How do these considerations relate to a subject's justification for believing propositions that are available to her as part of H? Recall that, according to J-Explanationism, the justification a subject has for believing p when p is available to her as part of an explanation H depends on the fact that that explanation can be expected, given her overall evidence, to represent her best attempt at accounting for e_H by virtue of containing the true proposition p or propositions that entail it. Now, to say that a subject can expect H to be the right way to account for e_H is simply to say that a subject is justified in believing the propositions that are available to her as part of H. This is because, a subject can expect H to be the right way to account for e_H if and only if that subject can expect, given her overall evidence, H to represent her best attempt at accounting for e_H . Hence, to be in a position to understand why e_H by means of H, a subject must, in addition to knowing how to account for e_H by means of H, be justified in believing the propositions that are available to her as being part of H.

In the case considered by Hills, assuming that the teacher's testimony truly defeats the justification the subject has for believing that Napoleon was a great leader, the subject cannot be regarded as being in a position to understand the facts about Napoleon reported in the book she read. Even if she maintains the conclusion that Napoleon was a great leader because of the way that explanation can account for him being tactically astute and so on, she cannot, given the teacher's testimony, expect that explanation to be the right way to account for those facts about Napoleon. From the perspective of her evidential situation, it is a matter of luck that the way in which the explanation she endorses can account for what she learned about Napoleon constitutes the right way to account for those facts. Accordingly, even if she maintains the conclusion that Napoleon was a great leader because of the way that explanation can account for him being tactically astute and so on, that conclusion cannot provide her with an understanding of what she learned about Napoleon. To gain that understanding, she must challenge her teacher's testimony on a rational basis and regain justification for maintaining the conclusion that Napoleon was a great leader.¹⁵

The requirement that a subject be in a position to expect that H is the right way to account for e_H can be seen as bridging the grasping requirement of understanding with its accuracy requirement. For a subject to be in a position to understand why e_H by means of H, H has to be the right way to account for e_H . But this accuracy requirement should not be conceived of as being completely external. It is not sufficient that the subject knows how to account for e_H by means of H and that H happens to be the right way to account for e_H . From the perspective of the subject's evidential situation, it

should not be a matter of luck that the way in which the explanation she endorses can account for e_H constitutes the right way to account for it.

3.6. Concluding Remarks

While, as just argued, being justified in believing the propositions that are available as part of an explanation of e_H is required for a subject to be in a position to understand why e_H by means of that explanation, a subject must, in addition, grasp the explanation in the required way. To elucidate the nature of that specific requirement, I examined the conditions under which a subject can come to grasp the structure instantiated by a particular explanation and argued that such a grasp is the result of an inferential process. By performing inferences on each part of the explanation, a subject can acquire a knowledge of the ways in which the elements cited in the explanation depend on each other. But, as I outlined, that inferential process does not simply allow the subject to acquire more knowledge about the explanation she comes to grasp. Gaining an inferential knowledge of the ways in which the elements of the explanatory story depend on each other amounts to acquiring a knowledge of how to account for the explanandum by means of the explanation.

The present chapter focused on the understanding a subject can gain of some portion of her evidence. Nevertheless, the import of my discussion of the grasping and justification requirements of understanding is not limited to the understanding a subject can gain of her own evidence. Even if a proposition is not part of a subject's evidence, the considerations put forward in the present chapter speak for the conclusion that gaining an understanding of why that proposition is true requires grasping an explanation and being justified in believing its content. Thus, while my present focus has been the understanding that can be gained by a subject of the constituents of her own evidence, the considerations put forward in the present chapter are relevant to a general elucidation of the notion of explanatory understanding.

Notes

- 1 By "explanatory understanding" I mean the type of understanding whose paradigmatic ascriptions involve the use of a why-clause such as in "S understands why mammals need oxygen" or "S understands why water solidifies at a certain temperature."
- 2 The precise connection between explanations and explanatory understanding is in fact a matter of controversy. Lipton (2009) offers several cases to cast doubt on the idea that a subject needs an explanation of e to understand why e is the case. On that particular point, I take Khalifa (2017) to be right in pointing out that if a subject understands why e , she should *at least* be able to offer an answer to the question "why e ?" which requires having some explanation of e at her disposal.
- 3 As already noted in Chapter 1, for a subject to be justified in believing a proposition p that is part of H given e , H does not need to constitute a potential explanation.

tion of every aspect of e . The fact that H could explain, if true, certain aspects of e is sufficient.

- 4 Here, I do not wish to restrict e_H 's content to propositions concerning a subject's own mental states. As noted by Austin "any kind of statement could state evidence for any other kind" (1962, p. 116) and, as I will argue in Chapter 5, there are good reasons to refrain from operating such a restriction.
- 5 See, *inter alia*, De Regt and Dieks (2005), De Regt (2009), Greco (2014), Grimm (2006, 2014), Pritchard (2010, 2014), Kelp (2015, 2017), Hills (2009, 2016), Strevens (2013, 2017), Khalifa (2017), Janvid (2018), and Dellsén (2020).
- 6 According to Grimm (2014), there are reasons to doubt that the son, in that case, even possesses the conceptual resources to genuinely believe the content of the explanation that is made available to him. I believe, however, that one can grant that the son possesses the required conceptual resources while still lacking a conception of how faulty wiring might cause a fire. What the son lacks, as I will argue, is the ability to make the appropriate conceptual connections between the elements of the account that is made available to him. But it does not follow that he lacks the concepts required to genuinely believe that the house burned down because of faulty wiring.
- 7 As Newman (2012) outlines, there are a large variety of inferences that subjects perform when trying to understand something by means of a theory. See Graesser, Singer, and Trabasso (1994), Cote, Goldman, and Saul (1998), and Graesser and Bertus (1998).
- 8 Jaworski (2009) distinguishes between three types of analytic how-questions: how-questions of means, how-questions of manner, and how-questions of mechanism. Here, my focus is on how-questions of mechanism although I do not want to rule out the possibility of other types of how-questions being relevant to the grasp a subject can gain of a particular explanation.
- 9 This is not to say that gaining a knowledge of the reasons why H explains e_H cannot allow S to grasp H in the required way. According to what I have claimed, gaining such knowledge is just not necessary to grasp H.
- 10 See also Dellsén (2020) for whom understanding essentially amounts to grasping models of dependence relations; the notion of grasp being understood, by Dellsén, as a shorthand for the relation between mind and models.
- 11 One might point out that performing inferences on each part of an explanation is not the only way to acquire such knowledge. A subject can be told that w is the way in which e_H depends on c . While this is entirely correct, I believe that such cases are most often cases where a subject is not merely provided with a particular piece of propositional knowledge but is guided through a conception of how the elements cited in a particular explanation relate to each other. Indeed, by stating that w is the way in which e_H depends on c , one can guide a subject into making certain inferences. Yet, in order to grasp how a particular explanation can account for e_H , the subject has to connect the elements cited in that explanation by drawing the relevant inferences herself.
- 12 While Grimm talks primarily about knowledge of causes, he (2014, p. 341) remains open to a conception of explanatory understanding that relies on a broader notion of dependence such as the one put forward by Kim (1994) which includes, but is not limited to, causal relations.
- 13 I take this conception of the grasping requirement of understanding to be radical in that it postulates the existence of a special ability to grasp modal reality directly.
- 14 Stanley and Williamson (2001) and Stanley (2011) offer a systematic analysis of know-how in terms of knowledge of propositions of the form ' w is the way to Φ '. This analysis of know-how fits particularly well with the proposed elucidation of the grasping requirement of understanding as it establishes a systematic

link between descriptions of the ways in which something can be done and knowledge of how to do it. However, it is important to note that the proposed elucidation of the grasping requirement of understanding involves a further step. According to it, it is by discovering the ways in which the elements of the explanatory story depend on each other, according to an explanation H, that a subject discovers how to use H to account for the explanandum. What the subject needs to discover is therefore not, primarily, how something can be done but the ways in which certain elements depend on each other according to a representation that could, if true, explain a certain phenomenon.

- 15 Dellsén (2018, 2019) offers another ground for the claim that understanding and justification come apart. He argues that for a subject to understand a given phenomenon by committing to a set of propositions P, P must be logically consistent and it must be rationally acceptable for her to commit to any logical consequence of P. Yet, according to Dellsén, constraints of logical consistency and deductive closure do not apply to justified belief. While I do not deny the first premise of Dellsén's argument, the considerations put forward in the previous chapters show why its second premise should be rejected, for, as I argued, constraints of logical consistency and deductive closure typically apply to justified belief.

References

- Austin, J. L. (1962). *Sense and Sensibilia*. Oxford: Oxford University Press.
- BonJour, L. (2005). In Defense of the a Priori. In M. Steup & E. Sosa (Eds.). *Contemporary Debates in Epistemology*. Malden: Blackwell, 98–105.
- Cote, N., Goldman, S. R. & Saul, E. U. (1998). Students Making Sense of Informational Text: Relations between Processing and Representation. *Discourse Processes*, 25, 1–53.
- Dellsén, F. (2017). Understanding Without Justification or Belief. *Ratio*, 30, 239–254.
- Dellsén, F. (2018). Deductive Cogency, Understanding, and Acceptance. *Synthese*, 195, 3121–3141.
- Dellsén, F. (2019). Rational Understanding: Toward a Probabilistic Epistemology of Acceptability. *Synthese*, 198, 2475–2494.
- Dellsén, F. (2020). Beyond Explanation: Understanding as Dependency Modelling. *The British Journal for the Philosophy of Science*, 71, 1261–1286.
- Elgin, C. (2017). *True Enough*. Cambridge: MIT Press.
- Friedman, M. (1974). Explanation and Scientific Understanding. *Journal of Philosophy*, 71, 5–19.
- Graesser, A. C. & Bertus, E. L. (1998). The Construction of Causal Inferences While Reading Expository Texts on Science and Technology. *Scientific Studies of Reading*, 2, 247–269.
- Graesser, A. C., Singer, M. & Trabasso, T. (1994). Constructing Inferences during Narrative Text Comprehension. *Psychological Review*, 101, 371–395.
- Greco, J. (2014). Episteme: Knowledge and Understanding. In K. Timpe & C. A. Boyd (Eds.). *Virtues and Their Vices*. Oxford: Oxford University Press, 285–302.
- Grimm, S. (2006). Is Understanding a Species of Knowledge? *The British Journal for the Philosophy of Science*, 57, 515–535.
- Grimm, S. (2014). Understanding as Knowledge of Causes. In A. Fairweather (Ed.). *Virtue Epistemology Naturalized: Bridges between Virtue Epistemology and Philosophy of Science*. Dordrecht: Springer, 329–345.

- Hills, A. (2009). Moral Testimony and Moral Epistemology. *Ethics*, 120, 94–127.
- Hills, A. (2016). Understanding Why. *Noûs*, 50, 661–688.
- Janvid, M. (2018). Getting a Grasp of the Grasping Involved in Understanding. *Acta Analytica*, 33, 371–383.
- Jaworski, W. (2009). The Logic of How-questions. *Synthese*, 166, 133–155.
- Kelp, C. (2015). Understanding Phenomena. *Synthese*, 192, 3799–3816.
- Kelp, C. (2017). Towards a Knowledge-Based Account of Understanding. In S. Grimm, C. Baumberger & S. Ammon (Eds.). *Explaining Understanding: New Essays in Epistemology and the Philosophy of Science*. London: Routledge, 251–271.
- Khalifa, K. (2017). *Understanding, Explanation, and Scientific Knowledge*. Cambridge: Cambridge University Press.
- Kim, J. (1994). Explanatory Knowledge and Metaphysical Dependence. *Philosophical Issues*, 5, 51–69.
- Kitcher, P. (1981). Explanatory Unification. *Philosophy of Science*, 48, 507–531.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lipton, P. (2009). Understanding Without Explanation. In H. W. de Regt, S. Leonelli & K. Eigner (Eds.). *Scientific Understanding*. Pittsburgh: University of Pittsburgh Press, 43–63.
- Newman, M. (2012). An Inferential Model of Scientific Understanding. *International Studies in the Philosophy of Science*, 26, 1–26.
- Newman, M. (2013). Refining the Inferential Model of Scientific Understanding. *International Studies in the Philosophy of Science*, 27, 173–197.
- Newman, M. (2017). Theoretical Understanding in Science. *British Journal for the Philosophy of Science*, 68, 571–595.
- Pritchard, D. (2010). Knowledge and Understanding. In A. Millar, D. H. Pritchard & A. Haddock (Eds.). *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press, 3–91.
- Pritchard, D. (2014). Knowledge and Understanding. In A. Fairweather (Ed.). *Virtue Epistemology Naturalized: Bridges between Virtue Epistemology and Philosophy of Science*. Dordrecht: Springer, 315–327.
- de Regt, H. (2009). Understanding and Scientific Explanation. In H. de Regt, S. Leonelli & K. Eiger (Eds.). *Scientific Understanding: Philosophical Perspectives*. Pittsburgh: University of Pittsburgh Press, 21–42.
- de Regt, H. & Dieks, D. (2005). A Contextual Approach to Scientific Understanding. *Synthese*, 144, 137–170.
- Stanley, J. (2011). *Know How*. Oxford: Oxford University Press.
- Stanley, J. & Williamson, T. (2001). Knowing How. *The Journal of Philosophy*, 98, 414–444.
- Strevens, M. (2013). No Understanding Without Explanation. *Studies in History and Philosophy of Science Part A*, 44, 510–515.
- Strevens, M. (2017). How Idealizations Provide Understanding. In S. Grimm, C. Baumberger & S. Ammon (Eds.). *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. New York: Routledge, 37–49.
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.

4 The Centrality of Truth

4.1. *The Accuracy Requirement*

At the end of the previous chapter, I noted that the justification requirement of understanding bridges its grasping requirement and its accuracy requirement. But what does this accuracy requirement precisely amount to? A subject who believes in a mythological explanation of the occurrence of thunder does not understand why thunder occurs even if she grasps that explanation in the required way and can expect it to be the right way to account for the occurrence of thunder. Likewise, a subject who believes that her house burned down because a witch cast a spell on her cannot be regarded as having an understanding of why her house burned down. This even if she grasps that explanation in the required way and, oddly enough, can expect it to be the right way to account for the fire. The reason for this is that the explanation by means of which that subject accounts for the fact that her house burned down is simply not accurate. Given her knowledge of how to account for the fact that her house burned down by means of that explanation, she may feel as if she understood why her house burned down. Yet, it is an entirely different question as to whether she *genuinely* understands why that occurred.

What does it take for an explanation H of e_H to be accurate in the relevant sense? A natural thought is that H is accurate whenever it constitutes a correct explanation of e_H , where a correct explanation of e_H is a set of true propositions that actually explain why e_H . There are, however, reasons to think that such a conception of the accuracy requirement of understanding is overly simplistic. Consider the following version of a case discussed by Elgin (2007):

Battle of Marathon: Sam claims to understand why the Athenians won the battle against the Persians at Marathon. A historian who is specialized in Ancient Greek history starts questioning him in order to find out whether he genuinely understands why the Athenians won that battle. Most of the answers given by Sam are correct but after several rounds of questions, he mistakenly claims that the Athenians deployed 10,000 troops at Marathon while, in fact, they deployed 9000 troops.

Sam's mistaken claim concerning the number of troops deployed at Marathon by the Athenians can be taken to show that the explanation of the Athenian victory over Persia he endorses is not the correct one as Sam states, falsely, that the Athenians deployed 10,000 troops at Marathon. Nonetheless, if all the other answers Sam gives to the historian are correct, finding out that he holds that false belief cannot, on its own, provide evidence that he does not genuinely understand why the Athenians won the battle against the Persians at Marathon. Intuitively, Sam's understanding of the Athenian victory at Marathon does not hinge on that false belief. To make this point even clearer, suppose that instead of falsely claiming that the Athenians deployed 10,000 troops at Marathon, Sam falsely claimed that they deployed 9001 troops. Holding that false belief surely wouldn't prevent him from understanding why the Athenians won against the Persians at Marathon.

This shows that understanding does not require complete accuracy. It is possible for a subject to understand why e_H while holding some false beliefs about e_H 's explanatory story. Yet, it is also true that holding certain false beliefs disqualifies one from counting as understanding why something is the case. Suppose for instance that, as it turns out, Sam believes that the Athenians used gunpowder during the battle of Marathon. Holding that particular belief surely prevents Sam from counting as having a genuine understanding of why the Athenians won against the Persians at Marathon.¹

4.2. *Central Truths*

Why is believing certain falsehoods about e_H 's explanatory story compatible with having an understanding of why e_H whereas believing other falsehoods is not? According to Kvanvig: "when the falsehoods are peripheral, we can ascribe understanding based on the rest of the information grasped that is true and contains no falsehoods" (2003, p. 201).² To account for the kind of contrast illustrated by the Battle of Marathon case, Kvanvig thus proposes to draw a distinction between propositions that are central to the understanding of a given phenomenon and propositions that are only peripheral to its understanding. The question is: what does it take for a proposition p to be central to the understanding of why e_H ?

Presumably, a proposition p that is central to the understanding of why e_H is a true proposition concerning e_H 's explanatory story such that being wrong about p —*i.e.*, about whether p is true—would prevent one from gaining an understanding of why e_H .³ In contrast, a proposition p that is only peripheral to the understanding of why e_H is a true proposition concerning e_H 's explanatory story such that one can be wrong about p and still qualify as understanding why e_H . With respect to the Battle of Marathon case, it is plausible that the proposition 'the Athenians did not use gunpowder during the battle of Marathon' is central to the understanding of the Athenian victory over Persia. For its part, the proposition 'the Athenians deployed 10,000 troops rather than 9000' is only peripheral to the understanding

of that historical event. It is still unclear, however, how precisely the distinction between propositions that are central to the understanding of a phenomenon and propositions that are only peripheral to its understanding should be construed. What makes the proposition ‘the Athenians did not use gunpower during the battle of Marathon’ central to the understanding of the Athenian victory? Without a clear answer to this question, the distinction introduced by Kvanvig does not tell us much more than we already know. Namely that believing certain falsehoods about e_H ’s explanatory story is compatible with having an understanding of why e_H while believing other falsehoods is incompatible with having such an understanding.

Intuitively, whenever a subject is wrong concerning a proposition that is central to the understanding of why e_H , that subject can be viewed as having an (apparent) understanding of why e_H that importantly differs, with respect to its content, from the (genuine) understanding one would have if one was correct regarding that proposition.⁴ In contrast, if a subject is wrong concerning a proposition that is only peripheral to the understanding of why e_H , it is reasonable to think that her understanding of why e_H is close, when it comes to its content, to the understanding one would have if one was correct concerning that proposition. For instance, if Sam’s only mistake regarding the battle of Marathon is to believe that the Athenians deployed 10,000 troops at Marathon while, in fact, they deployed 9000 troops, it is plausible that the content of his understanding is not importantly different from the content of the understanding one would have if one was not mistaken regarding that particular issue. In contrast, if Sam mistakenly believes that the Athenians used gunpowder during the battle of Marathon, the content of the (apparent) understanding he has of that historical event appears to be importantly different from the content of the (genuine) understanding one would have if one did not make that mistake.

The distinction introduced by Kvanvig thus appears to track a distinction that pertains to the content of the understanding (apparent or genuine) one has of a phenomenon. Moreover, it is plausible that the difference in content between the (apparent) understanding a subject has of why e_H by means of H and the understanding one would have by means of a correct explanation of e_H can be traced to differences in the ways each explanation accounts for e_H . Suppose that H_i is a correct explanation of e_H and that a subject accounts for e_H by means of another explanation H_j . If the way H_i explains why e_H differs importantly from the way H_j would explain it, then the content of the (apparent) understanding the subject has by means of H_j differs importantly from the content of the genuine understanding H_i can provide. This, I believe, accounts for the intuitions elicited by the Battle of Marathon case. An explanation of the Athenian victory stating that the Athenians used gunpowder would explain that historical event in a way that differs importantly from the way a correct explanation of the Athenian victory actually explains it. In contrast, an explanation stating that the Athenians deployed 10,000 troops would not explain that event in a way that is importantly

different from the way a correct explanation of the event explains it. As a result, if Sam's only mistake regarding the battle of Marathon is to believe that the Athenians deployed 10,000 troops at Marathon while, in fact, they deployed 9000 troops, the content of his understanding is not importantly different from the content of the understanding one would have if one was not mistaken regarding that issue. But if Sam mistakenly believes that the Athenians used gunpowder during the battle of Marathon, the content of the (apparent) understanding he has of that event is importantly different from the content of the (genuine) understanding one would have if one did not make that mistake.

This suggests that the distinction between propositions that are central and propositions that are only peripheral to the understanding of why e_H can be articulated in terms of the way in which an explanation explains or would explain why e_H is the case. Propositions central to the understanding of why e_H are propositions that make a difference when it comes to the way H explains why e_H in the following sense:

Centrality: A proposition p which is part of a correct explanation H of e_H is central to the understanding of why e_H iff had H stated that q rather than p , where q consists of either $\neg p$ or a proposition that entails $\neg p$, H would have explained why e_H in a way that importantly differs from the way H actually explains why e_H .⁵

This account manages to capture the intuition that if a subject is wrong concerning a proposition that is central to the understanding of why e_H , the content of her (apparent) understanding of why e_H importantly differs from the content of the understanding one would have if one was correct concerning why e_H is the case. One might worry, however, that Centrality is in fact sensitive to small changes that should not be factored into the results this account delivers. Suppose for instance that Sam falsely believes of the Athenian victory at Marathon that the Athenians deployed 8000 troops and suppose that deploying 8000 troops instead of 9000 would have required enacting an entirely different battle plan. In such a case, the way the account of the battle endorsed by Sam would explain the Athenian victory differs importantly from the way a correct account explains that victory. Consequently, Centrality delivers the result that the proposition 'the Athenians deployed more than 8000 troops' is central to the understanding of the Athenian victory. Now, such a consequence may seem problematic as what appears as a small mistake on Sam's part prevents him, given Centrality, from gaining an understanding of why the Athenians won the battle at Marathon. But, in my view, this consequence rather highlights a strength of Centrality. Indeed, if deploying 8000 troops at the battle of Marathon would have required enacting an entirely different battle plan, then believing falsely that the Athenians deployed 8000 troops is not a small mistake. If Sam mistakenly believes that the Athenians deployed only 8000

troops, he either fails to be sensitive to the way the Athenian battle plan was dependent on the number of troops they deployed or has a conception of that battle plan which differs importantly from the battle plan that was actually enacted by the Athenians. In both cases, the content of the (apparent) understanding Sam has of the Athenian victory by means of the account he endorses differs importantly from the content of the understanding one would have if one was correct concerning the number of troops deployed by the Athenians. Centrality manages to capture this aspect of the situation under examination although Sam's false belief appears, at a superficial level, to constitute only a small mistake on Sam's part.

It is worth noting that Centrality is also in line with the intuitive idea put forward by Gordon (2021, p. 4965) that propositions central to the understanding of why e_H are highly practically relevant propositions. Propositions central to the understanding of why e_H are taken by Gordon to be such that if a subject were to believe falsely whether those propositions are true, it would undermine her reliability when making predictions or manipulating information pertaining to e_H . In the previous chapter, I argued that a subject's ability to make predictions or manipulate information pertaining to e_H depends on her grasp of the explanation of e_H she endorses. Now, suppose that a subject grasps an explanation H of e_H and, as a result, is able to make predictions or manipulate information pertaining to e_H . If the way H would explain why e_H differs importantly from the way a correct explanation of e_H actually explains it, then the abilities the subject has by virtue of her grasp of H must differ from the abilities one would have by virtue of one's grasp of a correct explanation of e_H . The reason for this is that what the subject comes to know when grasping H is closely connected to the way H explains or would explain why e_H is the case.

4.3. Felicitous Falsehoods and Acceptance

Even though Centrality constitutes a promising account of the distinction that can be drawn between propositions central to the understanding of a particular phenomenon and propositions that are not, it raises certain questions when it comes to the role of idealizing assumptions in the understanding that can be gained of phenomena. There is no denying that scientific and everyday understanding depends on such assumptions. A well-known example concerns the ideal gas law. This law ($PV=nRT$) describes in a very simple way the relation between pressure (P), volume (V), and temperature (T) in a gas (n being the amount of substance and R being a constant). It also relies on several idealizing assumptions as it ignores the size of the gas molecules and the intermolecular attraction that takes place within gases. As a result, this law cannot be taken as defining the behavior of real gases but instead of ideal gases.

How can someone come to understand the behavior of a real gas by means of the ideal gas law? For any particular gas x , the proposition

‘intermolecular attraction takes place within x ’ is presumably central to the understanding of x ’s behavior. Finding out that someone mistakenly believes that no intermolecular attraction takes place within a particular gas would typically provide evidence that one does not genuinely understand the behavior of that gas. Yet, as the ideal gas law depends on the assumption that no intermolecular attraction takes place within a gas, accounting for the behavior of a particular gas by means of this law involves relying—either tacitly or explicitly—on this false assumption. To address this problem, note that a scientist who relies on the ideal gas law to account for the behavior of a real gas typically does not believe, of the target gas, that no intermolecular attraction takes place within it. Learning how to account for the behavior of gases by means of this law involves learning that it cannot be taken as defining the behavior of real gases. Accordingly, the notion of belief appears to be ill-fitted to characterize the type of commitment a scientist has toward the idealizing assumptions which the ideal gas law depends on. As suggested by Elgin (2004, 2017), the notion of acceptance seems to be better suited for this purpose.

It is generally admitted that the attitude of acceptance can be adopted in a particular context toward a proposition whether or not the accepted proposition is believed to be true.⁶ Consider the following characterization of this attitude provided by Cohen:

To accept that p is to have or adopt a policy of deeming, positing, or postulating that p – *i.e.* of including that proposition or rule among one’s premises for deciding what to do or think in a particular context, whether or not one feels it to be true that p .

(1992, p. 4)

Here, Cohen makes it clear that one can accept a proposition in a given context of inquiry regardless of whether one believes that proposition to be true or “feels it to be true” as he puts it. One can accept, for instance, that no intermolecular attraction takes place within a gas even if one knows, of the target gas, that intermolecular attraction takes place within it. This is because one’s acceptance of a particular proposition is motivated by prudential reasons which have no direct bearing on the truth of the accepted proposition. Typically, the reason why one accepts that no intermolecular attraction takes place within a gas in a particular context of inquiry is that this assumption allows describing in a very simple way the relation between pressure, volume, and temperature in that gas.

Now, if the type of commitment a subject has toward idealizing assumptions is best characterized in terms of contextual acceptance, should the understanding the subject has of the behavior of a gas by means of the ideal gas law not also be construed in terms of contextual acceptance? More generally, should the propositional commitment involved in understanding not be conceived of in terms of acceptance? As pointed out by Elgin (2017,

p. 97), the falsehoods which scientists rely on are not mere causal antecedents of the accounts they use. They are central to those accounts and, as a result, the notion of acceptance seems crucial to characterize the understanding those accounts can deliver.⁷ For his part, Dellsén argues that:

Understanding something may merely involve treating certain propositions or theories as given in the context of explaining something, as opposed to being disposed to feel that the propositions or theories are true.

(2017, p. 248)

Conceiving of the propositional commitment involved in understanding solely in terms of acceptance, however, raises important problems. As pointed out by philosophers such as Bratman (1992) or Cohen (1992), one crucial feature of acceptance is that it is tied to a particular context of reasoning.⁸ Yet, it is hard to make sense of a subject's cognitive standing with respect to the explained phenomenon if understanding is conceived of solely in terms of claims that are contextually accepted by that subject. Suppose that a subject accepts the content of a particular explanation in a specific context of inquiry. Assuming that the subject's acceptance of those claims is sufficient for her to understand the target phenomenon in that context, how should that subject's cognitive standing be characterized outside that particular context of reasoning? Plausibly, if a subject understands a particular phenomenon by means of some account, she retains that understanding (provided that the standards of epistemic appraisal do not change) although she is not in the process of explaining that phenomenon. That is, she retains that understanding although she is no longer in the specific context which prompted her to accept the content of the account she used to explain the phenomenon. But, if the propositional commitment involved in understanding is conceived of solely in terms of contextual acceptance, it is not clear how the content of the subject's understanding should be characterized outside the context which prompted her to accept the claims she used to explain the phenomenon. For outside that particular context, those claims may no longer be accepted by the subject.

According to Lawler (2021) a distinction must be drawn between propositions that facilitate a subject's understanding of a particular phenomenon and propositions that are elements of the content of that understanding. Contextual assumptions that are made in a particular context to explain a phenomenon belong to the first category. They facilitate the understanding that is gained of a phenomenon. In contrast, it is reasonable to think that propositions that are *believed* to be true of the phenomenon as a result of the use of certain contextual assumptions are elements of the content of the understanding that is gained. This is because the type of propositional commitment a subject has toward the elements of the content of the understanding she gains is not tied to a particular context of inquiry. Thus, even

if contextual acceptances play an important role in the understanding a subject gains of certain phenomena, the notion of belief remains better suited to characterize the subject's commitment toward the propositions that are elements of the content of the understanding she gains.

If these considerations are correct, they show that the tension between the role of idealizing assumptions in the understanding delivered by certain theories and the account of centrality offered in the previous section is only apparent. For the notion of centrality, as it has been explicated, concerns the content of the understanding a subject has of why e_H and depends on the claim that if a (true) proposition p is central to the understanding of why e_H , a subject cannot understand why e_H while being incorrect, with respect to what she *believes*, concerning p .⁹

4.4. *Understanding in Contexts*

The account of centrality put forward in this chapter relies on the idea that the way in which a potential explanation would explain why e_H can differ *importantly* from the way a correct explanation of e_H actually explains it. As already noted, there is an intuitive sense in which an explanation stating that the Athenians used gunpowder at the battle of Marathon would explain the Athenian victory in a way that importantly differs from the way a correct explanation explains it. But when is that difference sufficiently important for some proposition p to count as being central to the understanding of a particular phenomenon?

I believe that the threshold at issue here is best conceived of contextually. This is because, as outlined by philosophers such as Wilkenfeld (2013, 2017) and Kelp (2015, 2017), the notion of understanding itself exhibits a sensitivity to the overall context of epistemic appraisal.¹⁰ Wilkenfeld notes for instance that:

The understanding of someone who got a 4 on her AP calculus test might clearly count as understanding when being evaluated for one job opportunity (e.g. high-school summer intern), clearly not count when being evaluated for another (e.g. professorship in an elite math department), and neither clearly count nor not count when being evaluated for yet a third (e.g. admission as a student to an elite math department).

(2013, p. 1007)

Our readiness to ascribe understanding to a subject seems to depend on contextual factors that go beyond the accuracy and the justificatory status of that subject's cognitive states. In the example considered by Wilkenfeld, our readiness to ascribe understanding hinges on the question as to whether we are evaluating the subject as a potential high-school summer intern or as a potential math professor. Now, the threshold involved in Centrality, instead

of constituting a problematic aspect of the view, allows accounting, at least partly, for that dimension of understanding.

Consider another example discussed by Elgin (2007, p. 37) of a second grader who believes that humans evolved from apes. The belief of the second grader is false as, in fact, apes and humans evolved from a common hominid ancestor. Yet, in spite of her belief being false, one can legitimately attribute an understanding of human evolution to the second grader in a way that is not merely honorific.¹¹ The reason for this is that, plausibly, in a context in which the subject is being evaluated as a second grader, the proposition ‘humans and apes evolved from a common hominid ancestor’ is not central to the understanding of how humans evolved while the true proposition ‘humans evolved from another species,’ which I assume is also believed by the second grader, is. In such a context, the explanation stating that ‘humans evolved from apes’ does not deviate sufficiently from the way a correct explanation of how humans evolved actually proceeds for the proposition ‘humans and apes evolved from a common hominid ancestor’ to be central to the understanding of human evolution. In contrast, if a subject is being evaluated as an evolutionary biologist when it comes to her understanding of human evolution, finding out that she believes that humans evolved from apes would provide evidence for the conclusion that she does not genuinely understand how humans evolved.

It is thus reasonable to think that the threshold for some proposition to count as central to the understanding of why e_H depends on the overall context of epistemic appraisal. This does not mean, however, that the understanding of two subjects who are being evaluated in different contexts of appraisal cannot be compared. There is no denying that even if both the second grader and the evolutionary biologist can be attributed an outright understanding of human evolution, the second grader does not understand how humans evolved to the same extent as the evolutionary biologist. As understanding varies in degree, two subjects who are being evaluated in different contexts can be compared relative to the degree of understanding they have of the same phenomenon. Yet, how can such comparisons be made if what counts as central to the understanding of a particular phenomenon is sensitive to the overall context of epistemic appraisal?

Recall what I referred to as the whole truth in Chapter 2. According to what I claimed, the whole truth consists of an account that provides, *inter alia*, the best possible answer to the question “why is e_H the case?” Here, let me focus on the part of this account that is both necessary and sufficient to provide that answer—call it the whole truth about e_H . In addition, note that such an account can provide, if certain conditions are met, what Khalifa (2013a, 2017) and Kelp (2015, 2017) characterize as an ideal understanding of why e_H . Now, the degree of understanding a particular subject has of why e_H by means of an explanation H in a given context can be assessed in relation to the understanding that the whole truth about e_H would provide. If H explains e_H in the same way as the whole truth about e_H explains it,

then the degree of understanding that H can provide is ideal. In contrast, if the way in which H explains e_H differs from the way the whole truth about e_H explains it, then the degree of understanding H can provide is suboptimal. The degree of understanding that can be provided by a particular explanation can, therefore, be expected to depend on the extent to which the way that explanation would explain e_H deviates from the way the whole truth about e_H actually explains it.¹²

The question as to whether the way in which an explanation would explain e_H differs sufficiently from the way a correct explanation of e_H explains it thus depends on the overall context of epistemic appraisal. But the absolute degree of understanding that can be provided by an explanation depends on the extent to which the way that explanation would explain e_H deviates from the way the whole truth about e_H actually explains e_H . This means, essentially, that the threshold for outright understanding depends on the amount of propositions that count, in a particular context of epistemic appraisal, as central to the understanding of why e_H . In highly demanding contexts where many propositions are central to the understanding of why e_H , the degree of understanding one must have to count as outright understanding why e_H is higher than the degree one must have in less demanding contexts of epistemic appraisal.

4.5. *Understanding and Knowing*

Now that a precise elucidation of the accuracy requirement of understanding has been provided, let me come back more generally to the conditions that have to be met for a subject to be in a position to understand why e_H by means of an explanation H. According to the considerations put forward in the previous sections, H must allow accounting for e_H by means of the propositions that are central to the understanding of e_H in the context in which the subject is being evaluated. In addition, as outlined in the previous chapter, the subject must know how to account for e_H by means of H and she must be in a position to expect H to be the right way to account for e_H given her overall evidence. This yields, in light of what I argued concerning the way a subject's grasp of H is manifested in her understanding of why e_H , the following view of explanatory understanding:

Explanatory Understanding: S understands why e_H by means of H in C iff:

- (i) The propositions that are available to S as being part of H include the propositions that are central to the understanding of why e_H in C, and
- (ii) S believes with justification the propositions that are central to the understanding of why e_H in C because of the way H can account for e_H .

As it should be clear from the considerations put forward in Chapter 2, conditions (i) and (ii) can be satisfied without the subject knowing the propositions that are central in C to the understanding of why e_H . Yet, one may worry that the proposed view of understanding is too weak, for there is an intuitive sense in which understanding why e_H involves knowing why e_H is the case. There are, nonetheless, reasons for thinking that understanding's and knowledge's epistemic profiles differ substantially. In particular, understanding, contrary to propositional knowledge, seems to tolerate certain forms of epistemic luck.

Knowledge is widely regarded as being incompatible with epistemic luck and a good deal of the post-Gettier literature is premised on the idea that a subject cannot qualify as knowing that p , although she believes the true proposition p with justification, if her belief is true as a matter of luck.¹³ According to philosophers such as Kvanvig (2003) and Pritchard (2008, 2009, 2010), however, the same does not hold for understanding.¹⁴ Consider this slightly modified version of a case Pritchard takes to support this conclusion:

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

As Nero could easily have been wrong concerning why his house is burning had he formed his belief in the same way—that is, by asking what caused the fire to someone dressed as a firefighter—it seems that he is not in a position to know why his house is burning. But in Pritchard's (2010, pp. 78–79) view, it does not follow that he lacks an understanding of why his house is burning. The reason is that, according to Pritchard, the type of epistemic luck involved in this case is compatible with understanding.

One can draw a distinction between the kind of luck involved in standard Gettier cases such as the one discussed in Chapter 2 or Chisholm's (1966) sheep case—labeled intervening luck—and the kind of luck involved in cases such as Goldman's (1976) barn façade case—labeled environmental luck.¹⁵ In a case such as Chisholm's where a subject looking at a dog which perfectly occludes a real sheep forms the luckily true belief that there is a sheep in the pasture, luck “intervenes” between the belief and the fact. The subject is not really looking at a sheep in such a case. In contrast, in a case such as Goldman's where a subject forms the luckily true belief that there is a barn in the field she is looking at, luck is due to the epistemically inhospitable nature of the environment. It is due to the fact that the subject is driving through a region filled with fake barns.

While knowledge is vulnerable to both intervening and environmental luck, Pritchard argues that understanding is only vulnerable to intervening luck. In a case such as Nero and the Firefighters:

The agent concerned has all the true beliefs required for understanding why his house burned down, and also acquired this understanding in the right fashion. It is thus hard to see why the mere presence of environmental epistemic luck should deprive the agent of understanding.
(Pritchard, 2010, p. 79)

If Pritchard is correct and understanding, contrary to knowledge, is not vulnerable to environmental epistemic luck, then understanding why e_H is the case cannot consist of knowing why it is the case. Understanding why e_H simply does not require knowing the propositions that are central to the understanding of why e_H .

There are, however, reasons outlined by Khalifa (2013b, 2017) and Kelp (2015, 2017) for thinking that the case offered by Pritchard leads us astray as it does not really involve the type of luck present in a case like Goldman's. Kelp (2017, pp. 261–262) notes that one important disanalogy between the barn façade case and Nero and the Firefighter is that in the latter case, Nero would not have formed the same beliefs had he consulted a partygoer. He would have formed another set of beliefs concerning the explanation of why his house is burning that would have turned out to be false. In the barn façade case, though, the subject would have formed the same but false belief had she been looking at a fake barn. As Khalifa (2013b, pp. 8–10) outlines, in Nero and the Firefighter, luck seems to depend crucially on the fact that Nero acquired a particular bit of testimonial evidence as opposed to another one. In that, Nero is lucky to acquire non-misleading testimonial evidence. But this type of luck, as Pritchard (2005) himself remarks, is epistemically benign and does not undermine Nero's claim to know why the house is burning.

Boyd (2020), who agrees with Khalifa and Kelp regarding the weaknesses of the case initially considered by Pritchard, offers a variation which better matches the structure of a genuine Gettier case:

Claudius and the Electrician: On his way home, Claudius finds that all the houses in his neighborhood in Rome are on fire. Wondering what caused the fires, he stops by one of the houses, which happens to be Nero's. When Claudius asks a nearby electrician what caused the fire, she gives him the correct answer that it was a faulty breaker box. The electrician is trustworthy and knowledgeable, and there are no masquerading electricians nearby. On the basis of his true belief that the breaker box was faulty, plus background beliefs regarding electricity, the nature of breaker boxes, heat, and fires, Claudius forms a mental representation of the relationship between Nero's breaker box

and the fire: without a functioning breaker box to control the flow of electricity, wires in Nero's house were heated to the point that the surrounding wood combusted. And, in fact, Nero's house did burn down as a result of the faulty breaker box in the way that Claudius represents. Unbeknownst to Claudius, however, due to a recent outbreak of house fires, Roman officials decided that all houses would be retrofitted with lower-voltage breaker boxes. As a result, every house in the city has been retrofitted, except for Nero's, which still uses the old breaker box technology. Despite this change, rampant house fires persist as the result of faulty breaker boxes. When a faulty breaker box causes a fire in a retrofitted house, however, it is not for the same reasons that caused Nero's house fire: instead, the new breaker boxes give off sparks which ignite flammable materials that Romans like to store in their basements.

(Boyd, 2020, p. 83)

According to Boyd, Claudius and the Electrician truly allows addressing the question as to whether understanding tolerates environmental luck as, in this case, luck does not depend on the fact that Claudius obtained a particular bit of testimonial evidence as opposed to another one. The electrician's testimony provides Claudius with evidence for believing that without a functioning breaker box to control the flow of electricity, wires in Nero's house were heated to the point that the surrounding wood combusted—call that explanation H. As Claudius ignores that all the houses have been retrofitted except for Nero's house, he can rationally expect H to be the right way to account for the fire and as there are no masquerading electricians nearby, there is no question of evidential luck in that case. What makes Claudius's beliefs concerning why Nero's house burned down lucky is that, *given the nature of the environment*, it would have been easy for him to be wrong had he formed his beliefs in the same way. Had Claudius stopped in front of another house in his neighborhood, he would have falsely believed, of that house, that the fire was due to the fact that wires were heated to the point that the surrounding wood combusted.¹⁶ Consequently, it is reasonable to think that Claudius's true beliefs concerning why Nero's house burned down fall short of constituting knowledge.

With this in mind, the question is whether the luck present in Claudius and the Electrician prevents Claudius from understanding why Nero's house burned down. In Boyd's view, Claudius also falls short of understanding why Nero's house burned down as he fails to recognize that the fire in Nero's house constitutes an exception and is, therefore, unable to explain relevantly similar phenomena such as fires in neighboring houses. To quote Boyd directly:

if a hallmark of understanding is that it requires that one be able to not only provide a correct explanation in the relevant case, but in all cases

that one takes to be relevantly similar, then Claudius does not possess the relevant understanding in the individual case.

(2020, p. 85)

I see, however, reasons to resist Boyd's assessment of the case. While Boyd is right in claiming that Claudius is not able to explain correctly why the other houses burned down, one should pay close attention to the reason why he is unable to do so. The reason is that the cases Claudius takes to be relevantly similar to Nero's house are not, as a matter of fact, relevantly similar to that case. Unbeknownst to Claudius, all the other houses in the city were retrofitted and, as a result, he mistakenly takes other fires in the city to be similar to Nero's house. Yet, Claudius is perfectly able to offer a correct explanation in cases that are actually relevantly similar to Nero's house and it is not clear that his inability to correctly explain why the other houses burned down is due to a deficiency in his understanding of the individual case.

Consider the example of the diffraction patterns produced by crystals. Crystals are material which, when exposed to electromagnetic waves produce sharp or discrete diffraction patterns. Scientists understand this phenomenon based on the fact that the materials producing those diffraction patterns are three dimensional arrangements of atoms with translational periodicity—*i.e.*, they can be extended indefinitely in the three directions of space by repeating their unit cells. Yet, Shechtman, Blech, Gratias, and Cahn (1984) observed such diffraction patterns produced by three dimensional arrangements of atoms that lack translational periodicity. That discovery showed that the explanation by means of which the sharp diffraction patterns produced by certain materials was understood could not be applied to the materials observed by Shechtman et al. It did not show, however, that the scientists' understanding of the diffraction patterns produced by what are now called classical crystals as opposed to quasicrystals was defective. The translational periodicity of classical crystals is the reason why such material arrangements produce sharp diffraction patterns. What the discovery showed, essentially, is that the scope of the explanation by means of which scientists understood the sharp diffraction patterns produced by certain materials was more restricted than initially thought. Likewise, in Claudius and the Electrician, the scope of the explanation that Claudius comes to endorse is more limited than he takes it to be. But it does not follow that Claudius lacks an understanding of the individual case. What Claudius lacks is a specific piece of knowledge concerning the scope of the explanation by means of which he understands the individual case. The knowledge that the fire in Nero's house is exceptional with respect to its explanation.

One should distinguish between a subject's ability to provide a correct explanation in the cases that are actually relevantly similar to the one she understands and a subject's ability to discriminate between the cases that are relevantly similar to the one she understands and the cases that aren't. What Claudius lacks, because he ignores that every house in the city has

been retrofitted except for Nero's house, is the ability to properly discriminate between cases that are relevantly similar to Nero's house and cases that aren't. But his inability is not due to a deficiency in his understanding of why Nero's house is burning. As he knows how the fire in Nero's house depends on the elements cited in the correct explanation he comes to endorse, Claudius is perfectly able to provide a correct explanation in cases that are actually relevantly similar to the case of Nero's house. There is simply no reason to regard his ability to discriminate between cases that are relevantly similar to the case he understands and cases that are not as being essential to his understanding of the individual case.¹⁷

These considerations are, however, not sufficient to vindicate the claim that understanding tolerates environmental luck. As Khalifa (2017, p. 185) notes, two possible claims concerning the compatibility of understanding with environmental luck ought to be distinguished. First, that a subject can understand why e_H by means of H even if her beliefs concerning why e_H are true as a matter of environmental luck. Second, that two subjects S_1 and S_2 can understand why e_H to the *same degree* by means of H even if S_1 's beliefs—but not S_2 's—concerning why e_H are true as a matter of environmental luck. My discussion of Claudius and the Electrician pertains to the first of these two claims as it supports the conclusion that Claudius understands outright why Nero's house burned down in spite of the luck present in that case. But as understanding comes in degrees, one could argue, following Khalifa (2013b, 2017), that a subject cannot understand why Nero's house burned down to the fullest extent when such luck is present. The reason being, the argument goes, that explanatory knowledge represents the standard in relation to which lesser degrees of understanding are to be assessed. What must be done, therefore, is to determine whether two subjects S_1 and S_2 can understand why e_H to the same degree even if S_1 's beliefs—but not S_2 's—concerning why e_H are true because of environmental luck. Consider the following case due to Khalifa:

Fiona the Firefighter: Fiona is a firefighter who does an exemplary job in collecting all of the evidence that could be extracted from the embers of Nero's house. On the basis of this evidence, she comes to believe that a faulty breaker box is the root cause of the fire, and this is in fact true. However, at the moment of the fire, the breaker box malfunctioned at the same time that the grounding wire shorted. The two events were causally independent, and the shorted grounding wire did not actually cause the fire. But, had the shorted grounding wire (rather than the faulty breaker box) caused the fire, Fiona would have discovered exactly the same evidence and still believed that the breaker box's malfunctioning explains why Nero's house caught fire. By contrast, one of Fiona's colleagues, Carmen, considers both the breaker box and the grounding wire explanation. Moreover, she has all of Fiona's evidence, plus she takes some additional voltmeter readings at different points in

the house, and these readings allow her to eliminate the grounding wire explanation.

(2017, p. 196)

Given the evidence collected from the embers of Nero's house, it seems that Fiona can expect the faulty breaker box explanation to be the right way to account for the fire and, as she is an experienced firefighter, it can be assumed that she knows how the fire depends on the elements cited in that explanation. Yet, as the breaker box malfunctioned at the same time as the grounding wire shorted and the shorted wire could easily have caused the fire, Fiona's beliefs regarding why Nero's house burned down are true as a matter of luck. Due to the epistemically inhospitable nature of the environment, it would have been easy for her to be wrong concerning why Nero's house burned down had she formed her beliefs in the same way. Granted that this does not disqualify her from counting as understanding outright why Nero's house burned down, the question is: does she understand that event to the same degree as Carmen does?

Carmen, contrary to Fiona, considers the possibility that the fire was caused by a shorted grounding wire and, as a result, performs a test which allows her to rule out that possible explanation. Given her use of additional voltmeter readings, Carmen is not lucky in being correct concerning why Nero's house burned down. As Khalifa (2017, p. 198) puts it, her beliefs are safer than Fiona's because in nearby possible worlds in which Carmen falsely believes that the fire is due to a faulty breaker box, Carmen does not form that belief in the same way as in the actual world. In particular, what matters, according to Khalifa, is that Carmen's beliefs are safer than Fiona's as a result of a better scientific explanatory evaluation. Carmen considers a plausible alternative explanation of why Nero's house burned down and collects decisive evidence in favor of one of the possible explanations for the fire. For this reason, Khalifa argues, Carmen understands why Nero's house burned down better than Fiona. Her beliefs concerning why Nero's house burned down better approximate scientific explanatory knowledge than Fiona's.

One aspect of this case, however, deserves greater scrutiny. While it is true that Carmen's beliefs are safer than Fiona's, part of the problem appears to be that both Fiona and Carmen have reasons to regard the grounding wire explanation as a plausible explanation of the fire. As a result, both of them should try to rule out that alternative prior to accounting for the fire by means of the faulty breaker box explanation. To be sure, Carmen is not supposed to have any particular insight into the possible cause of the fire. Both are experienced firefighters with the same information concerning their environment. Accordingly, both have reasons to suspect that the fire is due to a shorted grounding wire. It follows that, in *Fiona the Firefighter*, it is not clear, irrespective of the luck present in the case, whether Fiona is even justified in endorsing the faulty breaker box explanation. After all, the grounding wire explanation, which she has reasons to regard as plausible,

is compatible with all the evidence she collected from the embers of Nero's house. All that this case might show is, therefore, that full understanding requires justification. But such a requirement, as outlined in Chapter 2, is weaker than requiring explanatory knowledge for full understanding. Consequently, to address the question as to whether the luck present in Fiona the Firefighter negatively impacts the understanding Fiona has, one ought to modify the case in such a way as to make it clear that both Fiona and Carmen are justified in endorsing the faulty breaker box explanation:

Fiona the lucky Firefighter: Everything is exactly the same apart from the fact that Roman officials recently decided that all the houses would be equipped with brand new grounding wires that cannot, due to the way they are designed, cause a short circuit. Fiona is aware of that as she personally helped equip some houses with the new type of grounding wires. Yet, unbeknownst to her, Nero's house is the only house in the city that was not equipped with new grounding wires and due to the damage caused by the fire, she cannot tell, through inspection alone, that the house is still equipped with old grounding wires. In contrast, Carmen, who arrives later at Nero's house, is aware that this house was not equipped with new grounding wires and, therefore, she takes additional voltmeter readings to determine if the fire was due to a shorted grounding wire.

Here, Fiona has no reason to regard the grounding wire explanation as a plausible alternative to the faulty breaker box explanation. Given her epistemic situation and the evidence she collected from the embers of Nero's house, she can rationally expect the faulty breaker box explanation to be the right way to account for the fire. Nevertheless, her beliefs concerning why Nero's house burned down remain unsafe as, due to the nature of the environment, she could easily have been wrong had she formed her beliefs in the same way. Yet it is no longer clear that Carmen understands why Nero's house burned down better than Fiona in that particular situation. In what sense is Fiona's understanding of what happened deficient compared to Carmen's understanding?

Regarding his original case, Khalifa (2017, p. 200) claims that Carmen's understanding is superior to Fiona's in at least four respects. First, Carmen grasps more coherence-making relations than Fiona as she grasps how the faulty breaker box and the grounding wire explanation relate to each other and to the voltmeter readings. Second, Carmen can answer more what-if-things-had-been-different questions than Fiona as she can predict how the voltmeter readings would have been different had the grounding wire caused the fire. Third, Carmen has more true beliefs than Fiona about the fire as, for instance, she correctly believes that even if the breaker box had not malfunctioned, the house would not have caught fire because of the grounding wire. Fourth, Carmen's beliefs represent a greater cognitive

achievement than Fiona's for Carmen has ruled out a particularly thorny alternative explanation of the fire.

Let me start with the first two reasons to regard Carmen as understanding why Nero's house burned down better than Fiona. As I argued, understanding involves a grasping requirement and if Khalifa is correct that Carmen grasps more coherence-making relations and is able to answer more what-if-things-had-been-different questions than Fiona, then there is ground to regard Carmen's understanding as being superior to Fiona's. Yet, I am doubtful that such ground actually exists. That Fiona has no reason to regard the grounding wire explanation as a plausible alternative to the faulty breaker box explanation does not entail that she lacks a grasp of how the two explanations relate to each other as well as to the voltmeter readings. As an experienced firefighter, Fiona can know that if Nero's house was equipped with old grounding wires, the fire could have been due to a shorted grounding wire. She can also know that voltmeter readings would have allowed determining whether the fire was due to a shorted grounding wire. The knowledge Fiona lacks about her environment does not prevent her from having a good grasp of how the two explanations relate to each other as well as to voltmeter readings. Having such a grasp simply does not depend on having an accurate representation of her environment. Similarly, the knowledge she lacks about her environment does not prevent her from being able to correctly answer counterfactual questions about the fire. Nothing prevents Fiona from reaching the conclusion that she would have had certain voltmeter readings if a grounding wire had caused the fire. Consequently, the fact that Carmen's beliefs concerning why Nero's house burned down result from ruling out the grounding wire explanation does not mean that her beliefs manifest a better grasp of the faulty breaker box explanation than Fiona's beliefs. Both Carmen and Fiona endorse the content of the faulty breaker box explanation because of the way that explanation could account for the evidence they collected in the embers of Nero's house and the knowledge Fiona lacks about her environment does not prevent her from grasping that explanation to the same extent as Carmen.

The only advantage Carmen has over Fiona is that, contrary to Fiona, she knows that Nero's house is still equipped with old grounding wires and that a grounding wire shorted prior to the fire. As such, it is true that she has more true beliefs pertaining to the fire than Fiona. However, it is not clear that those additional true beliefs allow her to understand what actually happened better than Fiona does. In Fiona the lucky Firefighter, Fiona identified in a perfectly reasonable way the reasons why Nero's house burned down. While she ignores that the fire could have been caused by a shorted grounding wire, this does not prevent her from understanding to the fullest extent what happened in that particular case. After all, she has no reason to suspect that the fire is due to a shorted grounding wire and, presumably, she has a very good grasp of what actually led to the fire. Simply acquiring the true belief that Nero's house is in fact equipped with old grounding wires and that a grounding wire shorted at the same time as the breaker box

faulted would not add anything to Fiona's understanding of what actually led to the fire.

As to the last reason provided by Khalifa in favor of the conclusion that Carmen's understanding is superior to Fiona's, it is a direct consequence of the claim he wants to defend based on the case under examination. It is true that Carmen's beliefs represent a greater cognitive achievement than Fiona's in the sense that Carmen ruled out a plausible alternative to the explanation she came to endorse. Yet this is a reason to believe that Carmen's understanding is better than Fiona's only under the assumption that understanding improves as it better approximates what Khalifa regards as scientific explanatory knowledge; an assumption that I think should be rejected in light of the above considerations.

Once properly examined, the considerations put forward by Khalifa in favor of the conclusion that Carmen's understanding is superior to Fiona's are thus unconvincing. In Fiona the lucky Firefighter, there is no reason to regard Fiona's understanding of what happened as being deficient compared to Carmen's understanding. Both Fiona and Carmen seem to have an equally good understanding of why Nero's house burned down although Fiona lacks explanatory knowledge. More generally, two subjects S_1 and S_2 can understand why e_H to the same degree by means of H even if S_1 's beliefs—but not S_2 's—concerning why e_H are true because of environmental luck. The fact that S_1 's beliefs are safer as the result of a scientific explanatory evaluation does not necessarily increase S_1 's understanding of why e_H . The reason for this is that what matters when it comes to understanding is not the strength of the connection between a subject's beliefs concerning why e_H and the fact that H is the correct way to account for e_H . Whenever H includes the propositions that are central to the understanding of why e_H and the subject, given her overall evidence, can expect H to be the right way to account for e_H , what matters is that the subject's beliefs concerning why e_H manifest a sufficient grasp of how e_H actually depends on the elements cited in H .

If my assessment of the cases considered in the present section is correct, then the proposed account of explanatory understanding is in no way too weak. Understanding simply does not require knowing the propositions that are central to the understanding of a given phenomenon. Before turning to different issues, let me, nevertheless, examine how the luck present in cases such as Chisholm's (1966) sheep case relate to understanding. Recall that, in Pritchard's (2009, 2010) view, understanding is only compatible with environmental luck as intervening luck undermines both knowledge and understanding. To support this claim, he (2010, p. 78) considers a modified version of Nero and the Firefighters in which the person who is asked by Nero why his house is burning is actually a partygoer dressed as a firefighter but, by sheer luck, provides Nero with the correct explanation of why his house is in flames. According to Pritchard, Nero cannot, in this version of the case, gain an understanding of why his house is burning by means

of the explanation that is made available to him, for “one cannot gain an understanding of why one’s house burnt down by consulting someone who, unbeknownst to you, is not a real fire officer but instead merely someone in fancy dress” (Pritchard, 2010, p. 78).

The alternative version of Nero and the Firefighters considered by Pritchard, however, suffers from the same weakness as the version discussed at the beginning of the present section. In both of them, the luck present seems to be due to the fact that Nero acquired a particular bit of testimonial evidence as opposed to another one. Let me, therefore, consider the general structure which, according to Khalifa (2017, p. 211), should be instantiated by a case involving genuine intervening luck with respect to understanding:

An agent is Gettier-lucky with respect to the belief that p explains why e_H if:

1. In the actual world,
 - a. the agent’s belief that p explains why e_H is true, and
 - b. the agent’s belief that p explains why e_H is produced by the fact that q explains why e_H (rather than the fact that p explains why e_H).
2. In a nearby possible world, everything is the same as the actual world except:
 - a. the agent’s belief that p explains why e_H is false.

In a case structured according to this model, luck is not due to the fact that the subject happens to acquire a particular bit of evidence as opposed to another one. It is due, instead, to the fact that some explanation is true in the actual world and that the subject endorses another explanation which also happens to be correct because of the former explanation. In the modified version of Nero and the Firefighters considered by Pritchard (2010), Nero’s belief that the fire is due to a faulty breaker box is not produced by the fact that another explanation of the fire happens to be correct but by the fact that a partygoer dressed as a firefighter happens, by sheer luck, to provide Nero with a correct explanation of why his house is in flames.

Now, given the model proposed by Khalifa, it is quite clear that if S_i ’s beliefs concerning why e_H are lucky in that sense, then S_i cannot understand why e_H to the same degree as S_j whose beliefs are not true as a matter of luck. The reason for this is that if S_i ’s beliefs fit this model, then the explanation by means of which S_i accounts for e_H is necessarily incomplete compared to the one endorsed by S_j . Indeed, the two explanations involved in the model are not, strictly speaking, alternative explanations, for both are correct explanations of e_H . It is assumed that both ‘ p explains why e_H ’ and ‘ q explains why e_H ’ are true and, as a result, S_i , who only believes that p explains e_H , only gets half of the story so to speak. It is, therefore, not surprising that, in the case of standard Gettier luck, the degree of understanding one has of why e_H increases once luck is removed. But let me stress

that, contrary to what Khalifa suggests, the strength of the epistemic access the subject has to what explains why e_H is not the crucial factor here. While cases in which S_1 's beliefs concerning why e_H are true as a matter of intervening luck are cases in which S_1 's understanding is suboptimal, such understanding should not be conceived of in relation to explanatory knowledge where explanatory knowledge would constitute the standard in relation to which lesser degrees of understanding are to be defined. For, as I argued, a subject can possess a full and complete understanding of why e_H although her beliefs concerning the propositions that are central to the understanding of why e_H fall short of constituting knowledge. The reason why intervening luck negatively impacts one's understanding of why e_H is thus not that, as argued by philosophers such as Grimm (2006) and Khalifa (2013b, 2017), understanding involves an anti-luck component. The reason is, instead, that situations in which intervening luck is present, when it comes to understanding, are situations in which the explanation one possesses for e_H is incomplete.

4.6. Concluding Remarks

According to what I have argued in the present chapter, understanding requires being correct concerning the propositions that are central to the understanding of the target phenomenon. Building on a distinction drawn by Kvanvig, I offered an account of what it takes for a particular truth to be central to the understanding of a phenomenon and explained to what extent the notion of centrality at issue is sensitive to the context of epistemic appraisal. The proposed elucidation of the accuracy requirement of understanding allowed formulating, in conjunction with the elucidation of its grasping component offered in the previous chapter, a general account of explanatory understanding which makes the connection between that cognitive standing and doxastic justification explicit. That account, however, raised the question as to whether explanatory understanding can be secured without knowledge of the propositions central to the understanding of a phenomenon. To address this particular question, I examined the connection between understanding and environmental luck and argued that, as understanding is compatible with this type of luck, it does not require knowledge of the propositions central to the understanding of a phenomenon. The epistemic profile of explanatory understanding is such that the notion of knowledge cannot be relied upon to elucidate the nature of explanatory understanding.

Notes

- 1 Note that I am assuming here that Sam grasps the explanation he endorses. Consequently, Sam's false belief that the Athenians used gunpowder during the battle of Marathon is tied to a conception of the way in which the Athenian victory depended on the use of gunpowder during the battle.

- 2 See also Carter and Gordon (2016).
- 3 I am assuming here that one cannot understand why e_H and suspend one's judgment regarding the truth of p if p is central to the understanding of why e_H . Accordingly, I take propositions that are central to the understanding of why e_H to be true propositions such that one must be correct concerning those propositions to understand why e_H .
- 4 The notion of apparent understanding I am relying on here is merely used as a way to draw a certain contrast and should not be taken as introducing a new kind of cognitive standing—*e.g.*, quasi-understanding. A subject has an apparent understanding of why e_H , in the sense I am using this expression here, whenever she accounts (with justification and grasp) for why e_H by means of H and fails to understand why e_H because of H 's inaccuracy.
- 5 Note that there can be multiple correct explanations of one particular phenomenon; some being more complete than others. As a result, Centrality should be understood as ranging over all the correct explanations of why e_H .
- 6 See van Fraassen (1984), Stalnaker (1987), Cohen (1992), Bratman (1992), and Cresto (2010).
- 7 To say that p is central to a particular account is not the same thing as saying that p is central to the understanding of a phenomenon in the sense specified in the previous section. Mizrahi (2012), for instance, argues that the idealizing assumptions which the ideal gas law depends on are not central to the understanding that is delivered by this law. They are mere causal antecedents of the law and can be successively removed to obtain Van der Waals law. The notion of centrality invoked by Mizrahi is importantly different from the one I have been concerned with here as it rather pertains to the dependence of the understanding that can be delivered by a given account on certain assumptions.
- 8 This feature of acceptance is also acknowledged by Dellsén (2017, 2018, 2019) who focuses on propositions that are treated as true by a subject in a particular context of explaining.
- 9 Note that the distinction between propositions that facilitate the understanding of a particular phenomenon and propositions that are part of the content of the understanding one gains of the phenomenon is orthogonal to the question as to whether and how falsehoods are epistemically valuable. Nothing I have said entails that the false propositions that facilitate the understanding of a particular phenomenon are not epistemically valuable. Certain idealizing assumptions seem to be ineliminable from our best scientific accounts. Batterman and Rice (2014) consider, for instance, idealized models that rely on thermodynamic limits which assume an infinite number of particles and an arbitrarily large volume in the target system. In those models, idealizing assumptions seem to be ineliminable because, without those assumptions, the models are no longer able to account for some relevant phenomena such as phase transitions. In addition, irrespective of the question as to whether idealizing assumptions are eliminable from our best scientific accounts, one has to explain how felicitous falsehood can facilitate one's understanding of phenomena. Elgin's (2004, 2007, 2017) notion of exemplification constitutes, in this regard, a promising way of accounting for the contribution of felicitous falsehood to understanding which is, in principle, compatible with what I have said concerning the accuracy requirement of understanding.
- 10 See also de Regt and Dieks (2005) and Bachmann (2020).
- 11 See Kvanvig (2003, 2009) for the claim that some uses of the term "understanding" are best interpreted as being honorific.
- 12 Note that the whole truth about e_H is not equivalent to any correct explanation of why e_H . A partial explanation of e_H can qualify as a correct explanation whereas it cannot consist of the whole truth about e_H as it does not provide the best possible answer to the question "why is e_H the case?"

- 13 In what follows I will work under the assumption that a subject's belief that p is true as a matter of luck in the relevant sense whenever she could easily have been wrong about the fact that p , had she formed the belief that p in the same way.
- 14 See also Morris (2012), Rohwer (2014), and Hills (2009, 2016).
- 15 See Pritchard (2005, 2010).
- 16 As the electrician simply states that the fire is due to a faulty breaker box, there is a nearby world in which Claudius stops in front of another burning house in his neighborhood, asks a nearby electrician what caused the fire and forms the same yet false belief that the fire is due to wires being heated to the point that the surrounding wood combusted.
- 17 Note that the same considerations apply to the Blacksmith case discussed by Grimm (2006, p. 521).

References

- Bachmann, M. (2020). The Epistemology of Understanding: A Contextualist Approach. *Kriterion*, 34, 75–98.
- Batterman, R. & Rice, C. (2014). Minimal Model Explanations. *Philosophy of Science*, 81, 349–376.
- Boyd, K. (2020). Environmental Luck and the Structure of Understanding. *Episteme*, 17, 73–87.
- Bratman, M. (1992). Practical Reasoning and Acceptance in a Context. *Mind*, 101, 1–16.
- Carter, A. & Gordon, E. (2016). Objectual Understanding, Factivity and Belief. In M. Grajner & P. Schmechtig (Eds.). *Epistemic Reasons, Norms and Goals*. Boston: De Gruyter, 423–442.
- Chisholm, R. (1966). *Theory of Knowledge, Foundations of Philosophy Series*. Englewood Cliffs: Prentice-Hall.
- Cohen, J. (1992). *An Essay on Belief and Acceptance*. New York: Clarendon Press.
- Cresto, E. (2010). Belief and Contextual Acceptance. *Synthese*, 177, 41–66.
- Dellsen, F. (2017). Understanding Without Justification or Belief. *Ratio*, 30, 239–254.
- Dellsén, F. (2018). Deductive Cogency, Understanding, and Acceptance. *Synthese*, 195, 3121–3141.
- Dellsén, F. (2019). Rational Understanding: Toward a Probabilistic Epistemology of Acceptability. *Synthese*, 198, 2475–2494.
- Elgin, C. (2004). True Enough. *Philosophical Issues*, 14, 113–131.
- Elgin, C. (2007). Understanding and the Facts. *Philosophical Studies*, 132, 33–42.
- Elgin, C. (2017). *True Enough*. Cambridge: The MIT Press.
- van Fraassen, B. (1984). Belief and the Will. *Journal of Philosophy*, 81, 235–256.
- Goldman, A. I. (1976). Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, 73, 771–791.
- Gordon, E. (2021). Elgin on Understanding: How Does It Involve KnowHow, Endorsement and Factivity? *Synthese*, 198, 4955–4972.
- Grimm, S. (2006). Is Understanding a Species of Knowledge? *The British Journal for the Philosophy of Science*, 57, 515–535.
- Hills, A. (2009). Moral Testimony and Moral Epistemology. *Ethics*, 120, 94–127.
- Hills, A. (2016). Understanding Why. *Noûs*, 50, 661–688.

- Kelp, C. (2015). Understanding Phenomena. *Synthese*, 192, 3799–3816.
- Kelp, C. (2017). Towards a Knowledge-based Account of Understanding. In S. Grimm, C. Baumberger & S. Ammon (Eds.). *Explaining Understanding: New Essays in Epistemology and the Philosophy of Science*. London: Routledge, 251–271.
- Khalifa, K. (2013a). The Role of Explanation in Understanding. *The British Journal for the Philosophy of Science*, 64, 161–187.
- Khalifa, K. (2013b). Understanding, Grasping, and Luck. *Episteme*, 10, 1–17.
- Khalifa, K. (2017). *Understanding, Explanation, and Scientific Knowledge*. Cambridge: Cambridge University Press.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Kvanvig, J. (2009). Responses to Critics. In A. Haddock, A. Millar & D. Pritchard (Eds.). *Epistemic Value*. Oxford: Oxford University Press, 339–352.
- Lawler, I. (2021). Scientific Understanding and Felicitous Legitimate Falsehoods. *Synthese*, 198, 6859–6887.
- Mizrahi, M. (2012). Idealizations and Scientific Understanding. *Philosophical Studies*, 160, 237–252.
- Morris, K. (2012). A Defense of Lucky Understanding. *The British Journal for the Philosophy of Science*, 63, 357–371.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2008). Knowing the Answer, Understanding and Epistemic Value. *Grazer Philosophische Studien*, 77, 325–339.
- Pritchard, D. (2009). Knowledge, Understanding and Epistemic Value. *Royal Institute of Philosophy Supplements*, 64, 19–43.
- Pritchard, D. (2010). Knowledge and Understanding. In A. Haddock, A. Millar & D. Pritchard (Eds.). *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press, 3–91.
- de Regt, H. & Dieks, D. (2005). A Contextual Approach to Scientific Understanding. *Synthese*, 144, 137–170.
- Rohwer, Y. (2014). Lucky Understanding Without Knowledge. *Synthese*, 191, 945–959.
- Shechtman, D., Blech, I., Gratias, D. & Cahn, J. W. (1984). Metallic Phase with Long-Range Orientational Order and No Translational Symmetry. *Physical Review Letters*, 53, 1951–1953.
- Stalnaker, R. (1987). *Inquiry*. Cambridge: Bradford.
- Wilkenfeld, D. (2013). Understanding as Representation Manipulability. *Synthese*, 190, 997–1016.
- Wilkenfeld, D. (2017). MUDDy understanding. *Synthese*, 194, 1273–1293.

5 On Knowledge

5.1. *From True Belief to Knowledge*

According to the account of evidential support put forward in the previous chapters, the evidence possessed by a subject can be adequately conceived of as what is in need of an explanation for that subject. The justification a subject has for believing claims whose content goes beyond what is manifest to her depends on the merits of the explanations of her evidence that are available to her. In addition, depending on her grasp of the explanations that are available to her, forming beliefs concerning those claims is what can allow the subject to gain an understanding of her evidence. The notion of evidential support as it has been elucidated in Chapters 1 and 2 is thus closely connected to the notion of understanding. As argued in the previous chapter, however, the reason for this is not that understanding requires knowing the propositions that are central to the understanding of a phenomenon, for understanding and propositional knowledge have importantly different epistemic profiles. Yet, as I will argue in the present chapter, this does not mean that the notion of explanatory understanding is not importantly connected to the notion of propositional knowledge.

What does it take for a subject to know something to be the case? As already outlined, knowing does not simply amount to having a justified true belief. Nonetheless, it is intuitively correct that the distinction between knowledge and mere true belief is to be traced back to the justification one has for holding a true belief. Consider the following question:

(Q₁) Why is S correct concerning the fact that p ?¹

Our willingness to attribute knowledge of p to a subject depends on the answer to be given to this question. Moreover, when raising a question such as (Q₁), we are interested in the answer to be given to the following question:

(Q₂) Why does S believe that p ?

Note, however, that when raising such a question we are not looking for any explanation of the subject's belief. We tend to be interested in answers that explain the subject's belief that p by showing why it is reasonable for her to hold that belief. That is, we tend to be interested in answers that explain why the subject believes that p by citing the evidence she has for holding that belief as opposed to citing what can be characterized as non-rational causes of her belief. Stating, for instance, that the subject believes that somebody crossed the field where she is currently standing because a specific neuronal pathway was activated in her brain can provide an explanation of why she believes that somebody crossed that field. But this answer does not explain her belief by showing why it is reasonable for her to believe that somebody crossed the field. In contrast, stating that the subject believes that somebody crossed the field because of the footsteps she discovered in that field explains why she believes that somebody crossed the field by showing why it is reasonable for her to hold that belief. This statement explains the subject's belief by citing evidence for the claim that somebody crossed the field.

The reason why we are interested in such answers to (Q_2) is, I submit, that the notion of knowledge fundamentally tracks situations in which the evidence possessed by a subject suffices to explain why she is correct concerning the fact that p and, in the present chapter, my intention is to examine the following view of knowledge:

K-Explanationism: S knows that p whenever S having e as evidence for p is sufficient to explain why S is correct concerning the fact that p .²

In the context of K-Explanationism, e does not necessarily stand for the subject's overall evidence. However, as I take the statement 'S believes that p because S has e as evidence for p ' to show why it is reasonable for the subject to hold that belief just in case this statement is correct and she is justified in believing that p , using e to represent a subset of her overall evidence in the context of K-Explanationism should not lead to any confusion.³ In addition, I take the statement 'S believes that p because S has e as evidence for p ' to be sufficient as an answer to (Q_1) whenever this statement can enable someone to *understand* why the subject is correct concerning the fact that p . As answers to why-questions typically aim at providing information that can promote one's understanding of some phenomenon, it is indeed reasonable to conceive of the conditions under which the statement 'S believes that p because S has e as evidence for p ' is sufficient as an answer to (Q_1) in terms of the understanding one can gain of why the subject is correct concerning the fact that p . As a result, K-Explanationism explicates knowledge in terms of the understanding that can be provided by an explanation of why a subject is correct concerning the fact that p and thereby establishes a close connection between the notion of knowledge and the notion of explanatory understanding.⁴

In light of these clarifications, one might worry that K-Explanationism, in fact, simply states that the notion of knowledge tracks situations in which a subject has the justified true belief that p and, thereby, conflicts with what was argued in Chapter 2. Indeed, a subject having e as evidence for p is sufficient to explain why she is correct concerning p only if the subject is correct concerning p and believes p because she has e . Unless these conditions are satisfied, one cannot genuinely come to understand why the subject is correct concerning p by means of the explanation ‘S has e as evidence for p .’ Nevertheless, things are more complicated. Recall Gettier’s example discussed in Chapter 2 in which Smith’s evidence supports believing that Jones owns a Ford and Smith infers from that proposition that either Jones owns a Ford or Brown is in Barcelona. As Jones does not own a Ford but, by sheer luck, Brown is in Barcelona, Smith ends up with the true justified belief that either Jones owns a Ford or Brown is in Barcelona. But citing the evidence Smith has for believing that Jones owns a Ford would not be sufficient to explain why Smith is correct concerning the fact that either Jones owns a Ford or Brown is in Barcelona. This explanation would be in some important sense misleading to an outsider as Jenkins (2006, pp. 144–145) puts it. An outsider can broadly be conceived of here as someone who does not have any special insight into Smith’s situation but who can, if provided the right explanation, come to understand why Smith is correct concerning the fact that either Jones owns a Ford or Brown is in Barcelona. More precisely, as an outsider, one neither has evidence that Smith lacks in favor of the truth of the proposition ‘either Jones owns a Ford or Brown is in Barcelona’ nor has access to facts such that if Smith had had knowledge of them, he would not have believed that proposition. Now, citing the evidence Smith has for believing that Jones owns a Ford to explain why Smith is correct concerning the fact that either Jones owns a Ford or Brown is in Barcelona would be misleading for someone who qualifies as an outsider. It would mislead her into assuming that Smith actually remembered or saw that Jones owns a Ford.

5.2. Evidence as *What Does the Explaining*

As just noted, citing the evidence Smith has for believing that Jones owns a Ford to explain why Smith is correct concerning the fact that either Jones owns a Ford or Brown is in Barcelona would mislead an outsider into assuming that Smith *remembered* or *saw* that Jones owns a Ford. The reason for this, plausibly, is that the evidence a subject has for believing p is sufficient to explain why she is correct concerning the fact that p when the subject having that evidence can be described in the explanation by success verbs such as “seeing” that something is the case or “remembering” that something is the case.

Suppose, for instance, that a subject believes that somebody crossed the field where she is currently standing. Stating that it visually appeared to her

that somebody was crossing the field could not be sufficient to explain why she is correct concerning the fact that somebody crossed the field. In contrast, stating that the subject saw somebody crossing the field would suffice to explain why she is correct concerning the fact that somebody crossed the field. One would not need any additional information to understand why the subject is correct concerning the fact that somebody crossed the field.

Let me therefore focus on explanations in which a subject's evidence for believing that p is cited using success verbs such as "seeing" or "remembering." What characterizes such explanations? While a subject's belief that p has to bear some causal relation to her seeing that p in order for that seeing to explain why she is correct concerning the fact that p , such explanations appear to function as arguments to the effect that the subject is correct concerning the fact that p . For instance, the statement 'S saw somebody crossing the field' does not appear to work as an explanation of why the subject is correct concerning the fact that somebody crossed the field merely by presenting one with the cause of her true belief. It works, instead, by presenting one with an argument to the effect that the subject is correct concerning the fact that somebody crossed the field. By citing the evidence the subject has for believing that somebody crossed the field, that statement shows that *it is to be expected*—to put it as Hempel (1965) does in his discussion of deductive-nomological explanations—that she is correct concerning the fact that somebody crossed the field.

The reason for this, according to philosophers such as Dretske (1969), Williamson (2000), Cassam (2009), and Stroud (2009), is that seeing that such and such is the case is a way of knowing that it is the case. By stating that a subject sees that the glass is broken, one does not merely state a cause of her true belief that the glass is broken but something that entails that she is correct concerning that fact. For seeing that p entails knowing that p is the case. Such a view of perceptual seeing is related to a distinction that is often drawn between propositional seeing whose paradigmatic ascription involves the use of a that-clause and objectual seeing which, contrary to propositional seeing, is not taken to have substantial epistemic implications. Seeing a glass that is broken is compatible with not being in a position to know that that glass is broken. This is because objectual seeing is not conceptual in the way propositional seeing is. A subject can, in principle, see a glass that is broken without possessing the concept of a glass. In contrast, seeing that the glass is broken seems to involve knowing that it is broken for such a state is factive and, as Cassam (2009, p. 585) puts it, plausibly involves recognizing that the glass is broken.

Yet, if propositional seeing is just a way of knowing that such and such is the case, focusing on situations in which an explanation can be given using success verbs such as "seeing" or "remembering" does not appear to be very promising when it comes to elucidating the notion of knowledge. There are, nonetheless, reasons to resist such a view of propositional seeing. In particular, cases have been provided by McDowell (2002), Turri (2010),

and Pritchard (2011a, 2012) to show that, contrary to the received view, propositional seeing is compatible with failing to believe what is seen to be the case:

Real Barn: Suppose, for example, that one is in a situation in which one is genuinely visually presented with a barn and circumstances are in fact epistemically good (there's no deception in play, one's faculties are functioning correctly, and so on). But now suppose further that one has been told, by an otherwise reliable informant, that one is presently being deceived (that one is in barn façade county, say), even though this is in fact not the case. Clearly, in such a case one ought not to believe the target proposition, and hence one cannot possibly know this proposition either. (Indeed, if one did continue to believe the target proposition even despite the presence of this undefeated defeater, then one would still lack knowledge.)

(Pritchard, 2012, p. 26)

In Pritchard's view, even if the subject does not believe that there is a barn because she ought not to believe it to be the case it does not follow that she does not see that there is a barn. As he (2012, pp. 26–27) notes, if the subject realizes that she has been misled, she will naturally take herself as having seen all along that there was a barn and, it seems, she will be perfectly right in doing so.

While I agree with Pritchard's conclusion that the subject will be right in taking herself as having seen all along that there was a barn in such circumstances, one crucial aspect of the case requires further examination. It is common to distinguish occurrent from dispositional beliefs. While occurrent beliefs can be conceived of as a type of mental assent—*i.e.*, as judgments—dispositional beliefs are best conceived of in terms of their dispositional profile which is characterized, at least partly, by a disposition to judge that the belief's content is true. Now, when it comes to propositional knowledge, the relevant notion of belief is the dispositional one as conceiving of knowledge solely in terms of occurrent belief would be far too restrictive. Many propositions that are known by a subject at a particular time are not judged to be true by her at that time. Yet, as Ranalli (2014) points out, cases such as the Real Barn case only seem to support the conclusion that propositional seeing and occurrent belief can come apart. Although the Real Barn case makes it plausible that the subject sees that there is a barn without occurrently believing that p , it is not clear, in such a situation, that she lacks the dispositional belief that there is a barn. In fact, according to Ranalli (2014, pp. 1244–1245), by way of seeing that there is a barn, the subject believes dispositionally that there is a barn and is thereby disposed to judge that this is the case. The subject's disposition to endorse the content of her dispositional belief is simply masked because of the undefeated defeater present in that case. But this, as just outlined, does not prevent her

from being right concerning the fact that there is a barn in a sense that is relevant to attribute that piece of knowledge to her.

To assess Ranalli's (2014) reading of cases such as the Real Barn case, let me return to the initial intuition in favor of the claim that propositional seeing entails knowledge of what is seen to be the case. As already noted, propositional seeing is conceptual in a way objectual seeing is not and part of the intuition supporting the claim under examination is that propositional seeing involves the deployment of the same capacities as the ones manifested in the formation of a belief. One should, however, be cautious in the identification of those capacities. While it is extremely plausible that seeing that p involves grasping that p , it is less clear that propositional seeing involves a recognition that what is seen is the case. As noted in Chapter 3 with respect to the grasp one can have of an explanation H , grasping a certain content should not be confused with endorsing that content. One can grasp H 's content while failing to endorse it and, similarly, it seems perfectly possible for someone to grasp that p without thereby endorsing or recognizing that p is the case.^{5,6} By grasping that p , one is able, as Elgin (2017) puts it, to wield that content to further one's epistemic ends which include, in the case of propositional seeing, acquiring knowledge by recognizing that what is seen to be the case is actually the case. Thus, while it is true that propositional seeing involves the deployment of the same capacities manifested in the formation of the belief that p , the recognitional aspect that is characteristic of belief—either dispositional or occurrent—can be separated from the capacities that are deployed in seeing that something is the case.

With this in mind, I submit that propositional seeing is, contrary to what Ranalli (2014) claims, best conceived of as involving a *disposition to believe* whatever is seen to be the case. Following Audi (1994), one can draw a distinction between dispositions to believe and dispositional beliefs.⁷ As Audi himself remarks, this distinction is often hard to draw when concrete cases are considered and might ultimately depend on empirical considerations. Nevertheless, such a distinction finds application when it comes to cases such as the Real Barn case. If propositional seeing involves knowing how to wield what is seen to further one's epistemic ends, it is reasonable to regard such an experience as disposing the subject to believe what she sees to be the case. This is because the most basic way in which what is seen can be used to further one's epistemic ends when one sees that p is by constituting the rational basis of one's recognition that p is the case. Now, according to this conception of propositional seeing and of the capacities that are deployed in that type of experience, it is true that, in the Real Barn case, the undefeated defeater masks a disposition that is characteristic of propositional seeing. But that disposition is a disposition to believe that there is a barn rather than a dispositional belief. As the undefeated defeater that is present prevents the subject from recognizing that there is a barn, there is no interesting sense in which the subject is correct concerning the fact that there is a barn.

Seeing that p thus does not involve being correct concerning the fact that p . By presenting one with factive evidence for p , propositional seeing disposes one to exploit an evidential situation that is such that exploiting it would involve being correct concerning the fact that p . Yet, the nature of the circumstances in which the subject finds herself can prevent that disposition from being manifested although she actually sees that p .

Before turning to another important issue related to the type of explanations under examination, let me consider another potentially problematic aspect of the Real Barn case. One may worry that the proposed reading of the case is in fact in tension with the considerations put forward in Chapters 1 and 2. How, given what I claimed concerning the relation of evidential support, can a subject's justification for believing that p —*i.e.*, that there is a barn—be defeated when the subject sees that p ? Let me first outline that in the case just considered, while I take p to be perceptually manifest to the subject, p is available to her as part of an explanation for her overall evidence. Consequently, although the case involves a subject seeing that p , it is importantly different from the case of introspective beliefs discussed in Chapter 1 as the type of situation it describes is the one over which J-Explanationism ranges. Now, according to what I claimed in Chapter 1, some piece of evidence e_i defeats the justification a subject has for believing that p when e_i is such that, when added to the subject's overall evidence, the explanation containing p no longer is the best explanation the subject has for her evidence. In the case just considered it is, however, doubtful that the other passenger's testimony can defeat the justification the subject has for believing that p in such a way. If p is perceptually manifest to the subject, how could an explanation stating that $\neg p$ be a better explanation of her overall evidence than an explanation stating that p ? But note that, given the version of J-Explanationism put forward in Chapter 2, there is another way for some piece of evidence e_i to defeat the justification a subject has for believing that p . Recall the Investigation case. In that case, the inductive evidence Sally has for believing that after completing the last steps of her investigative procedure, a new suspect that better fits the evidence she has gathered will emerge is not such that when added to Sally's overall evidence, the proposition 'Jeremy committed the burglary' is no longer part of the best explanation Sally has for her evidence. Instead, that evidence is such that when added to Sally's overall evidence, Sally can no longer expect the best explanation she has for her evidence to represent her best attempt at accounting for that evidence by virtue of containing the true proposition 'Jeremy committed the burglary.'

This, I believe, is precisely the way in which what the subject has been told in the Real Barn case defeats the justification she has for believing that there is a barn. Even if the fact that there is a barn is perceptually manifest to her, given what she is told, she can no longer expect the best explanation she has for her overall evidence to be better than the other explanations available to her by virtue of containing the true proposition 'there

is a barn.’ Even if, because it is perceptually manifest to her, the fact that there is a barn itself constitutes the subject’s propositional justification for believing that it is the case, that proposition is no longer probable enough conditional on her evidence for her to be justified in believing that there is a barn. As a matter of fact, there is no reason to assume that the evidence possessed by a subject is always certain for her. While Bayesian rules of conditionalization tend to encapsulate such an assumption, more general rules do not and are perfectly able to show how the probability of a particular claim conditional on uncertain evidence can be determined rationally.⁸ What happens in the Real Barn case is that the testimony of the reliable informant lowers the probability of the evidence the subject has for believing that there is a barn and, as a result, that subject can no longer expect the best explanation she has for her evidence to be better than the other explanations available to her by virtue of containing the true proposition ‘there is a barn.’

5.3. *Mere Appearances and Normalcy*

In the previous section, I relied on the assumption that evidence can be adequately described by the use of success verbs such as “seeing.” Yet, one might reject the idea that statements such as ‘S sees that *p*’ can work in the way just described to explain why a subject is correct concerning the fact that *p*. The main reason for this is the intuition that a subject’s evidence must consist of what is common between the epistemically good and epistemically bad cases. Take the example of a subject looking at a field that somebody is currently crossing. One might insist that the evidence the subject has for believing that someone is crossing the field cannot be adequately described by stating that she *sees* that someone is crossing the field. This is because the evidence possessed by the subject would be the same if she was hallucinating somebody crossing the field. According to this line of argument, her evidence should instead be characterized in terms of perceptual seemings or appearances. A characterization that would apply whether or not somebody was actually crossing the field.

Yet, as noted at the beginning of the previous section, statements such as ‘it visually appears to S that *p*,’ if understood as a characterization of a state whose content implies nothing about the way things are beyond appearances, cannot suffice to explain why a subject is correct concerning the fact that *p*. And, on closer examination, it is not even clear that they can explain why it is reasonable for the subject to believe that *p*. According to J-Explanationism, a subject’s visual appearance as of *p* can make it reasonable for her to believe that *p* insofar as *p* is part of the best explanation she has for that visual appearance. However, as Neta (2004, p. 309) outlines, such an explanation would involve an important gap as what appears to the subject in a certain way would be conceived of objectively—that is, as something that can appear differently from the way it is—while the explanandum

would be conceived of subjectively. Given this gap, how can p be part of the best explanation the subject has for her visual appearance as of p ?

Take a simple case of visual appearance where a particular object, say a pen, visually appears as of being blue to a subject. Intuitively, the proposition ‘the pen is blue’ is part of the best explanation the subject has for her visual appearance as long as she does not have any reason to suspect that the pen appears to her differently from the way it actually is. The subject has, it seems, a presumptive yet defeasible warrant for believing that the pen is blue. But is this explanation truly the best the subject has for her visual appearance? After all, the pen is conceived of, in the explanation itself, as something that can appear differently from the way it really is. It could be argued that the explanation containing the proposition ‘the pen is blue’ is simpler than the explanation containing, for instance, the proposition ‘a blue light is being projected on the pen which is in fact white.’ But in what sense, precisely, would the former explanation be simpler than the latter? It is not simpler in the sense that it relies on fewer postulates to explain the subject’s visual appearance. The former explanation relies on the postulate that the pen appears as it really is or that things, more generally, appear as they really are to explain the visual appearance. In contrast, the latter explanation does not rely on such a postulate. It cannot be argued that the postulate that things generally appear the way they are is itself made plausible by the subject’s overall evidence either. The subject’s overall evidence, if conceived of as constituted by states whose content implies nothing about the way things are beyond appearances, is neutral regarding this particular question.

The intuition that the proposition ‘the pen is blue’ is part of the best explanation the subject has for her visual appearance is, I believe, tied to the fact that, as outlined by Smith (2010, 2016), the situation in which it visually appears to her that the pen is blue and the pen is blue is explanatorily privileged over the situation in which it visually appears to her that the pen is blue and the pen is not blue. The former situation is, so to speak, the default or normal situation. Contrary to the latter situation, it does not require any special explanation in which some interfering factor—*e.g.*, hallucination-inducing drug—would be cited. Now, it is because the former situation is the normal one that a subject’s visual appearance as of the pen being blue can explain why it is reasonable for her to believe that the pen is blue. If that situation was explanatorily on a par with the situation in which it visually appears to the subject that the pen is blue and the pen is not blue, there would be no ground for regarding the proposition ‘the pen is blue’ as being part of the best explanation the subject has for her visual experience. Accordingly, Smith’s notion of normalcy allows accounting for the fact that the statement ‘it visually appears to S that the pen is blue’ suffices to explain why it is reasonable for a subject to believe that the pen is blue. It raises, however, the question as to why the situation in which it visually appears to a subject that the pen is blue and it is the case that the pen is

blue is the normal one. As Smith (2016, pp. 39–40) remarks, the relevant notion of normalcy is not statistical. The fact that the situation in which it visually appears to a subject that p and p is the case is the normal one does not depend on the fact that this situation is the most frequent one. Instead, I submit that the reason why that situation is the normal one is precisely because the subject being correct concerning the fact that p can be explained in the way I claimed it can; namely, by using success verbs such as “seeing” to describe the evidence possessed by the subject. It is because the subject’s experience can be conceived of as what can, in epistemically favorable circumstances, deliver her with evidence that suffices to explain why she is correct concerning the fact that p that the situation in which it visually appears to the subject that p and p is the case qualifies as the normal situation.⁹ If the subject’s experience could not be conceived of in such a way, there would be no reason to regard the situation in which it visually appears to her that p and p is the case as the normal one. This situation would be explanatorily on a par with the situation in which it visually appears to her that p and it is not the case that p as there would be no privileged connection between the experience undergone by the subject, conceived of subjectively, and her being correct concerning the fact that p . Hence the statement ‘it visually appears to S that p ,’ taken as a description of what is common between the epistemically good and epistemically bad cases, can explain why it is reasonable for a subject to believe that p . But what makes such an explanation possible is the fact that the subject’s experience can be adequately conceived of as what can, in epistemically favorable circumstances, deliver her with evidence that suffices to explain why she is correct concerning the fact that p . What makes such an explanation possible is the fact that, in epistemically favorable circumstances, the subject’s evidence can be described by the use of success verbs such as “seeing” when attempting to understand why she is correct concerning the fact that p .¹⁰

5.4. *Explanationism about Knowledge*

In the first section of this chapter, I claimed that K-Explanationism should not be understood as merely stating that the notion of knowledge tracks situations where a subject has the justified true belief that p . Nevertheless, given the considerations put forward in the previous sections, one might be tempted to understand this view as stating that the notion of knowledge tracks situations in which a subject believes that p on the strength of her *conclusive* evidence for p .¹¹ If the notion of conclusive evidence is understood as evidence that is such that it suffices to explain why a subject is correct concerning the fact that p , I see no particular problem with such a reading of K-Explanationism. However, given the way I conceive of the evidence a subject possesses when she sees or remembers that p , I believe that this view is in fact much more closely connected to previous explanationist accounts of knowledge offered by philosophers

such as Rieber (1998) and Jenkins (2006).¹² Those accounts are premised on the idea that whenever a subject knows that *p* a suitable connection between her belief that *p* and the fact that *p* exists and that the best way to articulate that connection is in terms of a particular type of explanation. Thus, according to Rieber, a subject knows that *p* whenever the fact that *p* explains why she believes that *p*. Likewise, for Jenkins (2006), a subject knows that *p* whenever *p* is a good explanation of the fact that she believes that *p* for someone who is not acquainted with the particular details of *S*'s situation.¹³ Although K-Explanationism focuses on explanations of why a subject truly believes that *p* rather than explanations of why she believes that *p*, this account is close to the ones put forward by Rieber and Jenkins. This is true, in particular, of situations where a subject sees or remembers that *p* for, in such situations, the subject has *p* itself as evidence for believing that *p*.

Indeed, I do not take success verbs such as “seeing” and “remembering” to describe a subject’s evidence in the sense of describing the state that constitutes that evidence. In my view, to state that a subject sees that *p* is to describe the way in which she has *p* itself as evidence for *p*. More precisely, to state that a subject sees that *p* is to describe the way in which *p* is manifest to that subject. Of course, one may resist this characterization of the evidence a subject has when she sees that *p* in light of the type of considerations put forward by Neta (2002, p. 670).¹⁴ One may argue that a subject’s evidence for believing that *p* cannot reasonably be regarded as being *p* itself, for it would make no sense for her to simply assert *p* when asked what her evidence for believing that *p* is. Let me, however, point out that a reasoning that is premised on *p* and has *p* itself as its conclusion is a perfectly valid one and that, therefore, the oddness outlined by Neta is not due to the fact that *p* itself cannot be used to reason to the conclusion that *p*. In fact, if, as I claimed in Chapter 1, evidence is what can be used as premise in good reasoning, that a subject can have *p* itself as evidence for believing that *p* should not be particularly surprising. The problem with asserting that *p* when asked for evidence in favor of *p* is not so much that a subject cannot have *p* itself as evidence for believing that *p*. It is that the subject’s assertion leaves the question as to whether she *has p* as evidence open. This is because simply asserting that *p* cannot answer the question as to what evidence the subject has. It merely expresses a belief in *p*. To answer this question, the subject must state that in the circumstances she is in, the conditions for having *p* as evidence are satisfied by stating, for instance, that she sees that *p*.

Thus, when a subject sees that *p* and believes that *p* on the basis of her factive support for *p*, what explains why she is correct concerning the fact that *p* is the fact that *p* itself. This is not to say, however, that K-Explanationism is equivalent to the accounts offered by Rieber (1998) and Jenkins (2006). To be sure, consider the following case which raises substantial difficulties for the type of account they offer:

Newton's Apple: An apple falls on Newton's head. By an amazing coincidence, Newton has a brain lesion such that, whatever rate the apple falls, the impact of the apple will (together with the lesion) cause Newton to believe that objects fall at precisely this rate. As it happens, the apple falls at 32ft./sec.², and Newton walks away with the belief that objects fall at this rate.

(Rieber, 1998, p. 201)

Here, the fact believed by Newton enters into the explanation of his belief but it is quite clear that Newton does not know that the apple falls at 32ft./sec.² This is because the fact that the apple falls at 32ft./sec.² does not explain Newton's belief for the right reason. More precisely, while this fact explains why Newton believes that the apple falls at 32ft./sec.², it does not explain it because Newton has this fact as evidence for his belief and this is precisely what prevents, according to K-Explanationism, Newton from qualifying as knowing that the apple falls at 32ft./sec.² That *p* explains a subject's belief that *p* is not sufficient for her to know that *p*. For *p* has to explain her belief for the right reason and it does when the subject has *p* itself as evidence for believing that *p*. K-Explanationism's focus on whether the evidence a subject has suffices to explain why she is correct concerning the fact that *p* thus makes a crucial difference. While in situations where *p* is manifest to a subject and she knows that *p* it is true that *p* itself enters into the explanation of her belief that *p*, what matters when it comes to the subject's knowledge that *p* is that *p* explains her belief for the right reason.

This aspect of K-Explanationism is also important when it comes to a problem analogous to the one discussed in Chapter 1 with respect to justification: factual knowledge about the future. It seems that we can have factual knowledge about the future and any account of knowledge should, at least in principle, allow for this possibility. But if *p* is a future state of affairs, it is not clear how it can enter into the explanation of something present—*i.e.*, the fact that a subject believes that *p*.¹⁵ While this problem is taken seriously by Goldman (1988) and Rieber (1998), Jenkins and Nolan (2008) consider that linguistic evidence strongly supports the conclusion that such backward explanations are possible.¹⁶ However, as pointed out by Byerly (2013), there are serious reasons to doubt that linguistic evidence supports this conclusion in the way Jenkins and Nolan claim it does. In addition, even if linguistic evidence actually supports this conclusion, stronger theoretical reasons may defeat it.¹⁷

For K-Explanationism, the problem of factual knowledge about the future does not arise in the same way. This is because, if *p* is a future state of affairs, it is implausible that a subject can have *p* itself as evidence for believing that *p*. In such a case, a subject cannot see or remember that *p*. But one may argue, as Neta (2002, p. 672) outlines, that present and past observations can *show* that *p* will be the case and that, as a result, present and past observations can suffice to explain why a subject is correct concerning the

fact that p will be the case.¹⁸ Accordingly, the problem of factual knowledge about the future, when it comes to K-Explanationism, does not concern the possibility of backward explanations. It concerns the possibility of having inductive knowledge based on evidence that shows that something will be the case and that thereby suffices to explain why one is correct concerning the fact that such and such will be the case.

5.5. Some Skeptical Problems

According to Rieber (1998), the main appeal of explanationist accounts of knowledge is their ability to make sense of skeptical paradoxes that are central to almost all attempts at providing a philosophical treatment of knowledge. The reason for this, he claims, is that such accounts support a specific contextualist reading of knowledge ascriptions.

Contextualism takes knowledge ascriptions to be dependent, with respect to their truth conditions, not only on the evidential situation of the subject to whom a piece of knowledge is ascribed but on the overall context of epistemic appraisal.¹⁹ This approach has been defended by philosophers such as Cohen (1988), DeRose (1995), and Lewis (1996) who tend to explain the context-sensitivity of knowledge ascriptions by specific rules governing them. But according to Rieber, one crucial advantage of explanationism about knowledge is that, given such a view, one does not have to postulate *ad hoc* rules governing knowledge ascriptions to explain why such ascriptions are sensitive to the overall context of epistemic appraisal. This is due to the fact that explanations are themselves context-sensitive. Indeed, explanations are taken by Rieber to be sensitive to relevant contrasts that are made salient in a given context of appraisal. Consider the following version of van Fraassen's (1980) Syphilis case. The statement "Smith had Syphilis" can, in principle, explain why Smith has got paresis. But if the case of Jones who also had Syphilis yet never got paresis was considered, the statement "Smith had Syphilis" would no longer be acceptable. What would need to be explained, once this contrast is made salient, is why Smith rather than Jones has got paresis.

Since Rieber takes a subject to know that p whenever p explains why she believes that p , he argues, in light of the context-sensitivity of explanations, that the question as to whether the subject knows that p can only be ascertained once the relevant explanatory contrasts are taken into account. Take Dretske's (1970) case where a subject forms the belief that there is a zebra in a cage she is looking at. The subject's belief is true as the animal in the cage is in fact a zebra. But suppose that she also happens to consider the possibility that the animal in the cage is not a zebra but a cleverly disguised mule. Or suppose that her skeptic friend tells her that the animal she is looking at could be, for all she knows, a cleverly disguised mule. According to Rieber (1998, pp. 196–197) the contemplation of such a possibility produces a shift in the context of epistemic appraisal. In the ordinary context—that is, in the

context where the possibility of the animal being a cleverly disguised mule is not made salient—the subject knows that there is a zebra in the cage as the fact that there is a zebra in the cage explains why she believes it to be the case. But in the skeptical context, the subject cannot know that there is a zebra in the cage. The reason for this is that, in that context, the question “why does the subject believe that there is a zebra in the cage?” is naturally understood as “why does the subject believe that there is a zebra rather than a cleverly disguised mule in the cage?” As the subject is, by stipulation, unable to tell the difference between a zebra and a cleverly disguised mule, Rieber (1998, p. 196) holds that in such a context the fact that there is a zebra in the cage cannot be regarded as explaining why the subject believes that there is a zebra in the cage.

While I do not contest the merits of Rieber’s contextualism over other contextualist approaches, I see two problems with his reading of Dretske’s case. First, Rieber assumes that a change in the explanatory target warrants the conclusion that the explanation we were primarily interested in no longer counts as an explanation. From the fact that p being the case does not explain why the subject believes that p rather than q , Rieber simply concludes that p being the case does not explain why the subject believes that p . Yet, as Lipton (1991) notes, the explanation of a contrast is not always an explanation of the fact alone. After all, in Dretske’s case, there is a clear sense in which, even in the context where the possibility of the animal being a cleverly disguised mule is made salient, the subject’s belief that there is a zebra in the cage is explained by the fact that there is a zebra in the cage.²⁰ Second, Rieber assumes that the reason why there being a zebra in the cage does not explain why the subject believes that there is a zebra in the cage in the skeptical context is that the subject is unable to tell the difference between a cleverly disguised mule and a zebra. However, this assumption, which is called into question by philosophers such as Pritchard (2012), seems to be importantly misleading when it comes to understanding what goes wrong with the subject’s belief.

As stated in the previous section, to ascertain whether the subject knows that there is a zebra in the cage, we are concerned with the question as to whether there being a zebra in the cage explains the subject’s belief for the right reason—*e.g.*, because she has that fact as evidence by virtue of seeing that there is a zebra in the cage. Now, considering the possibility of the animal being a cleverly disguised mule does not shift the context of appraisal in the way Rieber supposes it does but can show, by introducing a demand for explaining the contrast, that there being a zebra in the cage does not explain the subject’s belief for the right reason. The contrast that is made salient is, therefore, relevant to determine whether the subject knows that there is a zebra in the cage; though not in the way Rieber takes it to be.

As Rieber (1998, p. 196) himself notes, what explains the subject’s belief that there is a zebra rather than a cleverly disguised mule in the cage is not simply the fact there is a zebra in the cage. The subject holds that belief, in

part, because she has no particular reason to think that the animal in the cage may be a cleverly disguised mule. Now, this explanation of the subject's contrastive belief can show in what stating that she sees that there is a zebra in the cage to explain why she is correct concerning the fact that there is a zebra in the cage would be misleading for an outsider. It can show that it would mislead someone into falsely assuming that the evidence the subject has for believing that there is a zebra in the cage suffices to explain why she is correct concerning the fact that there is a zebra rather than a cleverly disguised mule in the cage. This is because, if the subject sees that there is a zebra in the cage, then she has that fact itself as evidence for believing that there is a zebra in the cage and does not need any justification independent from the one provided by her visual experience such as the testimony of an expert or a DNA test to determine whether it is true that the animal in the cage is a zebra.

Let me unpack the proposed reading of Dretske's case. As just mentioned, what explains the subject's contrastive belief is that she has no particular reason to think that the animal she is looking at may be a cleverly disguised mule. In other words, she has no reason to take that possibility seriously and to refrain from believing that there is a zebra in the cage if she lacks justification independent from her experience for the claim that the animal in the cage is not a cleverly disguised mule. But, of course, there might be such reasons. If, for instance, the zoo keepers regularly disguised mules as zebras, there would be a reason for the subject to take that possibility seriously and the explanation of her contrastive belief would constitute a description of an epistemic limitation on her part.²¹ And, it is precisely by describing such a limitation that the explanation of the subject's contrastive belief can show to what extent explaining why she is correct concerning the fact that there is a zebra in the cage by stating that she sees that there is a zebra in the cage would be misleading. Explaining why the subject is correct by stating that she sees that there is a zebra in the cage would lead an outsider into assuming that the subject does not need any justification independent from the one provided by her visual experience to determine whether the animal in the cage is a zebra. Yet, if the explanation describes such a limitation, there are reasons for the subject to refrain from believing that there is a zebra in the cage in case she lacks justification independent from her experience for the claim that the animal in the cage is not a cleverly disguised mule.²²

In contrast, if there are no such reasons, the evidence the subject has is sufficient to explain why she is correct concerning the fact that there is a zebra rather than a cleverly disguised mule in the cage. This is true even if the subject is not able, through inspection alone, to tell the difference between a zebra and a cleverly disguised mule. As Pritchard (2012) outlines, her experience, even if it does not result from the exercise of an ability to discriminate between zebras and cleverly disguised mules, can provide her with evidence that *favors* one alternative over the other in such a way as to put her in a position to know that the animal in the cage is a zebra rather than a cleverly disguised mule. After all, in the situation under examination, there is no

reason for the subject to refrain from believing that there is a zebra in the cage in case she lacks independent justification for the claim that the animal in the cage is not a cleverly disguised mule. The subject does not need any justification independent from the one provided by her visual experience to determine whether it is true that the animal in the cage is a zebra.

If this reading of Dretske's case is correct, K-Explanationism does not support a contextualist reading of knowledge ascriptions. This is because the contrast to be explained when the possibility of the animal being a cleverly disguised mule is made salient does not show that there being a zebra does not explain why the subject believes that there is a zebra in the cage. What it can show, depending on the reasons there are for the subject to take the possibility of the animal being a cleverly disguised mule seriously, is that there being a zebra in the cage does not explain her belief for the *right reason* and that, therefore, explaining why she is correct concerning that fact by stating that she sees that there is a zebra in the cage would be misleading. It would mislead one into assuming, upon consideration of the skeptical possibility, that the evidence the subject has is sufficient to explain why she is correct concerning the fact that there is a zebra rather than a cleverly disguised mule in the cage. Thus, contrasts that are made salient in the context of appraisal can reveal that p does not explain the subject's belief that p for the right reason but the reason why p explains the subject's belief is not itself dependent on the context of appraisal.

One might, nevertheless, point out that as understanding is itself context-sensitive, K-Explanationism is bound to have contextualist consequences as it explicates the notion of knowledge in terms of the understanding one can gain of the fact that a subject is correct concerning p . According to what I argued in the previous chapter, understanding is indeed context-sensitive as the propositions that are central to the understanding of a particular phenomenon may depend on the overall context of epistemic appraisal. Yet, note that I did not rule out the possibility of some propositions being central to the understanding of a particular phenomenon in every context of epistemic appraisal. Presumably, the proposition 'human beings were not created as they are now by an all-powerful being' is central, in every context of epistemic appraisal, to the understanding of human evolution. Likewise, if a subject is correct concerning p because she sees that p , it is reasonable to think that the proposition 'S sees that p ' is central to the understanding of why the subject is correct concerning p in every context of epistemic appraisal. If a subject believes that p on the basis of her seeing that p , one cannot, irrespective of the context of appraisal, come to understand why she is correct concerning p without getting that fact right. Consequently, while the centrality of certain propositions depends on the overall context of appraisal, I take the propositions that are relevant to whether a subject knows that p according to K-Explanationism to be central to the understanding of why the subject is correct concerning p in every context of appraisal.

Let me now come back to the issue of skepticism more generally. I have just claimed that if the explanation of a subject's contrastive belief does not describe an epistemic limitation on her part, then, irrespective of whether the subject can tell the difference between a cleverly disguised mule and a zebra, her evidence can be sufficient to explain why she is correct concerning the fact that there is a zebra rather than a cleverly disguised mule in the cage. Accordingly, a subject can be in a position to know that the animal she is looking at is not a cleverly disguised mule; a consequence that appears perfectly acceptable in the kind of cases just considered. Yet, is such a consequence still acceptable when radical skeptical hypotheses such as the ones introduced by Descartes (1641/1979) or Putnam (1981) are considered? Consider the following argument:

(P₁) S does not know that she is not a handless brain in a vat.

(P₂) If S does not know that she is not a handless brain in a vat, then S does not know that she has hands.

(C) S does not know that she has hands.

(P₁) is intuitively correct and insisting that the subject is in a position to know the denial of the skeptical hypothesis would beg the question. The evidence she has simply does not appear, irrespective of the circumstances in which she is, sufficient to explain why she is correct concerning the denial of such a hypothesis. But why is this the case? Let me offer the following answer: irrespective of the circumstances in which the subject is, there are reasons for her to take the possibility of being a handless brain in a vat seriously. This is because the brain in a vat hypothesis is such that there is, in principle, no way for her to tell the difference between being a handless brain in a vat and being an embodied subject connected to reality through a reliable perceptual apparatus. It is by virtue of its very nature that the brain in a vat hypothesis provides a reason for anyone who grasps its content to take this possibility seriously. Now, from the above considerations it follows that, irrespective of the circumstances which the subject is in, the evidence she has is not sufficient to explain why she is correct concerning the fact that she has hands rather than being a handless brain in a vat.

This naturally prompts the conclusion that the skeptical argument is essentially correct and that a subject is never in a position to know that she has hands. Yet, this conclusion follows only insofar as the reasons there are for the subject to take the possibility of being a brain in a vat seriously are understood as reasons not to believe that she has hands unless she has justification independent from her experience for the claim that she is not a handless brain in a vat. Understanding those reasons in such a way would, however, be a mistake when it comes to radical skeptical hypotheses. The reason for this is that, as Coliva (2015, p. 96) points out, the specificity of such hypotheses is that independent justification for their denial cannot

possibly be obtained. There is no DNA test or expert testimony that could, in addition to the evidence provided by the subject's own experiences, allow her to determine whether she is an embodied subject rather than a handless brain in a vat.

Here, I propose to follow Wright (1986, 2004) in operating a distinction between evidential and non-evidential entitlement in the following sense:²³ understanding why a subject is correct concerning the fact that she has hands does not involve conceiving of that subject as being entitled, given the evidence she has for believing that she has hands, to reject the skeptical possibility. It involves, instead, conceiving of her as being non-evidentially entitled to reject the skeptical possibility. Accordingly, what the skeptical argument shows is that our route to the understanding of a subject's epistemic access to facts is dependent on a particular conception of that subject. It depends on conceiving of her as being non-evidentially entitled to reject the radical skeptical possibility made explicit in the argument. This argument, therefore, derives its strength from the fact that it forces upon us a conception of the subject which is such that once we adopt it, we are no longer in a position to understand why she is correct concerning the fact that she has hands by means of the evidence she has for holding that belief. It forces us to look for the subject's evidential basis for rejecting the skeptical possibility and, thereby, to conceive of her as being in need of an evidential basis for rejecting that possibility. Yet, if I am correct, once the subject is conceived of in such a way, there is no longer a route to the understanding of why she is correct concerning the fact that she has hands by means of the evidence she has.²⁴ In fact, once that conception is forced upon us, there is no longer a route to seeing how the subject can have evidence for that belief in the first place. In the third section of this chapter, I claimed that an appearance as of p can support believing that p as long as the subject does not have any reason to suspect that things appear to her differently from the way they are. The reason for this is that the situation in which it appears to her that p and p is the case is the normal situation. But if, irrespective of the circumstances in which a subject is, there are reasons for her to take the possibility of being a handless brain in a vat seriously, then, irrespective of the circumstances, there are reasons for her to suspect that things do not appear as they are. And unless the subject is conceived of as being non-evidentially entitled to reject the radical skeptical possibility, there is no route—not even the one examined in the previous sections—to seeing how the subject can have evidence for believing that she has hands.

Now, if a subject's entitlement for rejecting the claim that she is not a handless brain in a vat is strictly non-evidential, she cannot, given K-Explanationism, be in a position to know that she is not a handless brain in a vat and (P_1) is true. The question then is whether (C) is also true. The considerations just put forward regarding the way the skeptical argument works strongly hints at a negative answer to this question. If what the argument essentially shows is that understanding why a subject is correct concerning the fact that she

has hands involves conceiving of her as being non-evidentially entitled to reject the possibility made explicit in the argument, then this argument is not sufficient to show, on its own, that (C) is true. This is because showing this, given the proposed reading of the argument, would involve showing that our understanding of why the subject is correct concerning the fact that she has hands is somewhat defective. Yet, as such, the skeptical argument does not show that our understanding of the subject's epistemic access to that fact is in any way defective. We can, as a matter of fact, understand why she is correct concerning the fact that she has hands by means of explanations such as 'S sees that she has hands' or 'S remembers that she has hands' and the skeptical argument does not provide any reason to revise that judgment. While radical skeptical arguments show something important concerning the nature of the understanding one can have of a subject's epistemic access to facts and, thereby, concerning the nature of knowledge, they do not show that one is unable to achieve that understanding.

The proposed reading of the radical skeptical argument thus supports accepting (P_1) and rejecting (C) and (P_2), which amounts to limiting in scope knowledge's closure under known entailment. If understanding why a subject is correct concerning the fact that she has hands involves conceiving of her as being non-evidentially entitled to reject the possibility of being a handless brain in a vat, one cannot possibly understand why the subject is correct concerning the denial of the skeptical hypothesis by means of the evidence that puts her in a position to know that she has hands. This calls, in turn, for a corresponding limitation of the closure principle discussed in Chapter 1. Recall that, according to what I claimed, J-Explanationism encapsulates the idea that if a subject is justified in believing that p given the evidence she possesses, then whenever q is available to her as a logical consequence of p , that subject has sufficient evidence for believing that q . As shown in Chapter 1, this principle is perfectly acceptable in the large majority of cases. However, when it comes to skeptical hypotheses such a principle cannot hold, for, as outlined above, once one conceives of a subject as being in need of an evidential basis for rejecting the possibility of being a handless brain in a vat, there is no longer a route to seeing how she can have evidence for the belief that she has hands in the first place. Understanding a subject's epistemic access to facts involves conceiving of her as being non-evidentially entitled to reject hypotheses such as the brain in a vat hypothesis and the very possibility for her to acquire justification for believing mundane propositions depends on the fact that she can adequately be conceived of in such a way.

5.6. Concluding Remarks

In Chapter 4, I argued that explanatory understanding is not to be explicated in terms of propositional knowledge, for gaining an explanatory understanding of a particular phenomenon does not require knowing the propositions that are central to its understanding. But, as shown in the present chapter, it

does not follow that the notion of explanatory understanding is not importantly connected to the notion of propositional knowledge. According to the view of knowledge put forward in the present chapter, propositional knowledge is to be explicated in terms of explanatory understanding. This notion fundamentally tracks situations in which the evidence possessed by a subject suffices for someone to gain an understanding of why that subject is correct concerning a particular fact.

My discussion of the types of situation in which citing the evidence possessed by a subject suffices to provide one with such an understanding led me to argue that a subject's evidence can be adequately described by the use of success verbs such as "seeing." Indeed, it is by characterizing the way in which certain facts are manifest to a subject that an explanation can provide one with the relevant understanding of why that subject is correct regarding a particular claim. I then turned to skeptical challenges and claimed that radical skeptical arguments, while not sufficient to establish the conclusion that we generally lack knowledge of reality, shed light on the nature of the understanding one can have of a subject's epistemic access to reality. In particular, they show that such an understanding depends on conceiving of the knowing subject as being non-evidentially entitled to reject hypotheses that are typically made explicit in skeptical arguments.

Notes

- 1 I take (Q₁) to be equivalent to the questions "how did S arrive at the truth regarding *p*?" and "why does S truly believe that *p*?"
- 2 Note that K-Explanationism is not intended as an analysis of the notion of knowledge. It aims, instead, at elucidating this notion by offering a characterization of the types of situation this notion tracks.
- 3 The scope of K-Explanationism is thus limited to situations in which the statement 'S believes that *p* because S has *e* as evidence for *p*' shows why it is reasonable for the subject to believe that *p*.
- 4 As I argued that explanatory understanding is not itself to be explicated in terms of propositional knowledge, K-Explanationism should not prompt any worry of circularity. Yet, one may point out that while I did not explicate explanatory understanding in terms of knowledge, I explicated the grasping requirement of explanatory understanding in terms of propositional knowledge. For, according to what I claimed in Chapter 3, to grasp an explanation H of e_H in the required way is to know how e_H depends on the elements cited in H. However, to explicate the grasping requirement of explanatory understanding in terms of a specific type of inferential knowledge is not to explicate explanatory understanding itself in terms of such knowledge. Understanding why e_H by means of H in C involves believing with justification the propositions that are central to the understanding of why e_H in C because of the way H can account for e_H . Now, of course, for a subject to believe those propositions because of the way H can account for e_H , she must grasp H in the required way. But the inferential knowledge in terms of which the subject's grasp of H is explicated does not constitute her understanding of why e_H in the way her knowledge of the propositions that are central to the understanding of e_H would according to a knowledge-based account of explanatory understanding. This inferential knowledge merely enables the subject to

- acquire an understanding of why e_H by allowing her to form beliefs concerning why e_H in the way required to understand why e_H . Consequently, the elucidation of the grasping requirement of explanatory understanding offered in Chapter 3 does not raise substantial problems for the view I intend to examine here.
- 5 Philosophers such as Gordon (2012) equate understanding that p with knowing that p . Here, while I do not wish to deny that propositional understanding is reducible to propositional knowledge, I draw a distinction between grasping that p and understanding that p . One obvious reason for this is that if believing that p requires grasping that p , grasping that p cannot simply be equated with knowing that p .
 - 6 Note that the notion of grasp at issue here is importantly different from the one discussed in Chapter 3. In Chapter 3, my main concern was the grasp a subject can have of the structure instantiated by a particular explanation. In the context of the present discussion, however, the notion of grasp is more directly related to a subject's ability to articulate conceptually a particular content in order to use it as a rational basis for making certain judgments.
 - 7 Note that Audi (1994) holds the view that seeing that p involves the dispositional belief that p and not merely a disposition to believe that p .
 - 8 Bayesian conditionalization prescribes one to conditionalize one's expectations according to the following rule where Pr_f stands for a proposition's final probability, Pr_i for a proposition's initial probability and $\text{Pr}_i(e)$ is assumed to be maximal: $\text{Pr}_f(p) = \text{Pr}_i(e | p) \cdot \text{Pr}_i(p) / \text{Pr}_i(e)$. Jeffrey's conditionalization, in contrast, prescribes one to conditionalize one's expectations according to the following rule which covers cases where e is certain but allows determining p 's final probability in light of uncertain evidence: $\text{Pr}_f(p) = \text{Pr}_i(p | e) \cdot \text{Pr}_i(e) + \text{Pr}_i(p | \neg e) \cdot \text{Pr}_i(\neg e)$.
 - 9 Smith's (2016, p. 42) own explanation of why a given situation is normal relies on the idea that possible worlds can be ordered with respect to their relative normalcy: the situation where both p and q are true being normal just in case q is true in all the most normal worlds in which p is true. Yet, as noted in Belkoniene (2019), this answer depends on assumptions concerning what makes a world more normal than another that appear problematic when certain cases involving purely statistical evidence are considered.
 - 10 Let me stress that the statement 'it visually appears to S that p ', while being sufficient to explain why it is reasonable for the subject to believe that p , is not sufficient to explain why she is correct concerning the fact that p . My claim here is that what makes that statement an adequate explanation of why it is reasonable for the subject to believe that p is the fact that, in favorable epistemic circumstances, the subject's evidence can be described by the use of success verbs.
 - 11 See Neta (2002) for such a view.
 - 12 K-Explanationism also bears important similarities with the virtue-theoretic accounts of knowledge proposed by Zagzebski (1996), Greco (2002), or Sosa (2007, 2011). A common trait of those accounts is that they take a subject to know that p whenever her true belief that p results from or is held because of the exercise of certain intellectual virtues or abilities. While K-Explanationism locates the explanation at the level of the evidence the subject has for believing that p rather than at the level of her intellectual abilities, this view nonetheless encapsulates central intuitions driving virtue-theoretic accounts of knowledge. For it also elucidates knowledge in terms of an explanation of why a subject truly believes that p that is relative to the exercise of central intellectual abilities: a subject's ability to exploit her evidential situation.
 - 13 See also Goldman (1984, 1988) for whom a subject knows that p whenever p enters prominently into the best explanation of the fact that she believes that p .
 - 14 In Neta (2018), Neta endorses the view that p itself can be part of a subject's evidence when she sees that p .

- 15 As noted by Rieber (1998, pp. 200–201) a similar problem arises in the case of mathematical knowledge.
- 16 In order to deal with knowledge of future facts, Goldman (1988, p. 36) complicates his account by distinguishing two types of explanations for a belief. Rieber (1998, p. 200), for his part, considers certain conceptions of causation, notably Lewis's (1986), as possible solutions to this problem.
- 17 The very possibility of backward causal explanations gives rise to certain paradoxes. See, for instance, Mellor (1991) who argues that, given such paradoxes, the possibility of backward causal explanations can be ruled out *a priori*.
- 18 More generally, whenever a subject's knowledge that p is inferential, the evidence that suffices to explain why she is correct concerning p is not constituted by p itself. Yet, there is no reason to assume that the evidence a subject has for believing that p is sufficient to explain why she is correct concerning the fact that p only if she has p itself as evidence for p . Situations in which the evidence a subject has entails that p and in which the subject correctly infers that p from the evidence she possesses are paradigmatic examples of situations where the evidence the subject has suffices to explain why she is correct concerning p without p itself being part of her evidence.
- 19 Neta's (2003) own brand of contextualism actually takes the subject's evidential situation itself to be dependent on the overall context of appraisal.
- 20 In his discussion of Goldman's (1976) barn façade case, Rieber (1998, p. 201) assumes, in order for his account to deliver the right result, that the subject's belief that there is a barn in the field is not explained by the fact that there is a barn in the field. But in that case, too, it seems clear that what explains the subject's belief is the fact that there is a barn in the field even if the possibility of it being a fake barn is made salient in the context of appraisal.
- 21 The reasons for the subject to take the possibility of the animal being a cleverly disguised mule seriously are facts such that, if possessed as evidence by the subject, they would defeat the justification she has for believing that there is a zebra in the cage. While misleading defeaters considered in the second section of this chapter also defeat the justification a subject has for believing a particular proposition, they ought to be distinguished from the type of facts under examination. Misleading defeaters typically defeat the justification a subject has for believing that p because, if possessed as evidence by her, they are evidence for a false proposition q such that, if q was true, q would constitute a reason for the subject to take the possibility of p being false seriously.
- 22 Similar considerations apply to cases involving environmental luck such as Goldman's (1976).
- 23 See Coliva (2015) and Pritchard (2015, 2016) for responses to radical skepticism similar to Wright's in that they are also inspired by Wittgenstein's (1969) so-called hinge epistemology.
- 24 Let me insist on the fact that, given this reading of the skeptical argument, what the subject is non-evidentially entitled to reject is not merely a matter of our own choosing. For the route to the understanding of her epistemic access to facts is not itself a matter of our own choosing. What the subject is non-evidentially entitled to reject is something that we discover when faced with skeptical arguments and not merely something that we arbitrarily posit for the need of a particular inquiry.

References

- Audi, R. (1994). Dispositional Beliefs and Dispositions to Believe. *Noûs*, 28, 419–434.
- Belkoniene, M. (2019). Normalcy, Understanding and the Problem of Statistical Evidence. *Theoria*, 85, 202–218.

- Byerly, R. (2013). Explanationism and Justified Beliefs about the Future. *Erkenntnis*, 78, 229–243.
- Cassam, Q. (2009). Knowing and Seeing: Responding to Stroud's Dilemma. *European Journal of Philosophy*, 17, 571–589.
- Cohen, S. (1988). How to Be a Fallibilist. *Philosophical Perspectives*, 2, 91–123.
- Coliva, A. (2015). *Extended Rationality: A Hinge Epistemology*. Basingstoke: Palgrave Macmillan.
- DeRose, K. (1995). Solving the Skeptical Problem. *Philosophical Review*, 104, 1–52.
- Descartes, R. (1641/1990). *Meditations on First Philosophy*. Notre Dame: University of Notre Dame Press.
- Dretske, F. (1969). *Seeing and Knowing*. Chicago: University of Chicago Press.
- Dretske, F. (1970). Epistemic Operators. *Journal of Philosophy*, 67, 1007–1023.
- Elgin, C. (2017). *True Enough*. Cambridge: MIT Press.
- van Fraassen, B. (1980). *The Scientific Image*. Oxford: The Clarendon Press.
- Goldman, A. H. (1984). An Explanatory Analysis of Knowledge. *American Philosophical Quarterly*, 21, 101–108.
- Goldman, A. H. (1988). *Empirical Knowledge*. Berkeley: University of California Press.
- Goldman, A. I. (1976). Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, 73, 771–791.
- Gordon, E. (2012). Is There Propositional Understanding? *Logos and Episteme*, 3, 181–192.
- Greco, J. (2002). Virtues in Epistemology. In P. Moser (Ed.). *Oxford Handbook of Epistemology*. New York: Oxford University Press, 287–315.
- Hempel, C. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.
- Jenkins, C. (2006). Knowledge and Explanation. *Canadian Journal of Philosophy*, 36, 137–163.
- Jenkins, C. & Nolan, D. (2008). Backwards Explanation. *Philosophical Studies*, 140, 103–115.
- Lewis, D. (1986). Causal Explanation. In D. Lewis (Ed.). *Philosophical Papers Vol. II*. New York: Oxford University Press, 214–240.
- Lewis, D. (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- Lipton, P. (1991). Contrastive Explanation and Causal Triangulation. *Philosophy of Science*, 58, 687–697.
- McDowell, J. (2002). Responses. In N. H. Smith (Ed.). *Reading McDowell on Mind and World*. London: Routledge, 269–305.
- Mellor, D. (1991). Causation and the Direction of Time. *Erkenntnis*, 35, 191–203.
- Neta, R. (2002). S Knows That P, *Nous*, 36, 663–681.
- Neta, R. (2003). Contextualism and the Problem of the External World. *Philosophy and Phenomenological Research*, 66, 1–31.
- Neta, R. (2004). Skepticism, Abductivism, and the Explanatory Gap. *Philosophical Issues*, 14, 296–325.
- Neta, R. (2018). Your Evidence Is the Set of Facts That Are Manifest to You. In V. Mitova (Ed.). *The Factive Turn in Epistemology*. Cambridge: Cambridge University Press, 32–49.
- Pritchard, D. (2011). Epistemological Disjunctivism and the Basis Problem. *Philosophical Issues*, 21, 434–455.

- Pritchard, D. (2012). *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Pritchard, D. (2015). *Epistemic Angst: Radical Scepticism and the Groundlessness of Our Believing*. Princeton: Princeton University Press.
- Pritchard, D. (2016). Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing. *Philosophy and Phenomenological Research*, 93, 70–90.
- Putnam, H. (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Ranalli, C. (2014). Luck, Propositional Perception, and the Entailment Thesis. *Synthese*, 191, 1223–1247.
- Rieber, S. (1998). Skepticism and Contrastive Explanation. *Nous*, 32, 189–204.
- Smith, M. (2010). What Else Justification Could Be. *Noûs*, 44, 10–31.
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford: Oxford University Press.
- Sosa, E. (2011). *Knowing Full Well*. Princeton: Princeton University Press.
- Stroud, B. (2009). Scepticism and the Senses. *European Journal of Philosophy*, 17, 559–570.
- Turri, J. (2010). Does Perceiving Entail Knowing? *Theoria*, 76, 197–206.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wittgenstein, L. (1969). *On Certainty* (G. E. M. Anscombe & G. H. von Wright, Eds.). D. Paul & G. E. M. Anscombe (Tr.). Oxford: Blackwell.
- Wright, C. (1986). Facts and Certainty. In *Proceedings of the British Academy, Volume 71*. Oxford: Oxford University Press, 429–472.
- Wright, C. (2004). Warrant for Nothing (and Foundations for Free)? *Aristotelian Society Supplementary Volume*, 78, 167–212.
- Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.

6 The Possession of Truth

6.1. *From Facts to Evidence*

In my discussion of K-Explanationism, I argued that unless a subject's experience can be conceived of as what can, in favorable epistemic circumstances, provide the subject with evidence that suffices to explain why she is correct concerning certain claims about the world, one cannot account for the fact that experiences, conceived of subjectively, contribute toward the justification a subject has for believing such claims. This means, essentially, that a subject's evidence cannot be limited to facts about certain non-factive mental states she is in as certain philosophers would have it.¹

According to proponents of explanationist approaches of evidential support such as Conee and Feldman (2008) and McCain (2014a), explanationism broadly construed is particularly well suited to show how certain non-factive mental states can themselves contribute toward the justification a subject has for believing claims whose content goes beyond what is introspectively manifest to her. Recall, for instance, that in Conee and Feldman's view: "perceptual experiences can contribute toward the justification of propositions about the world when the propositions are part of the best explanation of those experiences that is available to the person" (2008, p. 97). Here I do not wish to deny that facts about a subject's experiences can contribute toward the justification she has for believing certain claims about the world. In my discussion of the justification one can have for holding introspective beliefs such as 'I have a tickle sensation,' I noted that when a subject attends to the tickle sensation she is undergoing, the fact that she has that sensation is manifest to her and, therefore, is part of her evidence. Conee and Feldman are right in pointing out that this fact can support believing propositions whose content goes beyond mere appearances. The proposition 'Bob is tickling me,' for instance, can be part of the best explanation the subject has for the fact that she is undergoing a tickle sensation. But if what I claimed in the previous chapter is correct, external facts that are seen or remembered by a subject are part of her evidence too.² Both internal and external facts can contribute toward the justification a subject has for believing

certain claims about the world. What distinguishes these two types of facts, as evidence, is what they put a subject in a position to *know* of the world. Facts about a subject's experiences, conceived of subjectively, cannot put her in a position to know propositions whose content goes beyond mere appearances. This is because such evidence cannot suffice to explain why she is correct concerning certain claims about the world. In contrast, external facts that are perceptually or memorially manifest to a subject can suffice to explain why she is correct concerning such claims. Those facts can be used to gain an understanding of the subject's epistemic access to the external world.

Note, however, that my discussion of the type of situations in which a subject qualifies as knowing certain claims to be true focused exclusively on the use of the subject's evidence that can be made by someone else to gain a particular understanding of the subject's epistemic access to facts. Indeed, I explicated the notion of knowledge in terms of the understanding *one* can have of why a subject is correct concerning *p* by means of the evidence possessed by the subject and, to gain such an understanding, one presumably needs independent evidence for thinking that the subject sees or remembers that such and such is the case. After all, that understanding is no different, in kind, from the understanding the subject herself can gain of what constitutes her evidence. Plausibly, it also requires having justification for endorsing a particular explanation of why the subject is correct concerning *p*. Yet, if the facts that can allow one to understand why a subject is correct concerning *p* can be adequately conceived of as being perceptually or memorially manifest to the subject, the access the subject has to those facts should be expected to be substantially different from the access an outsider who understands why the subject is correct concerning *p* has to those facts.

6.2. *Reflectively Accessible Evidence*

Part of the appeal of the view that a subject's evidence is limited to facts about certain of her non-factive mental states comes from the thought that for something to serve as a subject's evidence, the subject must enjoy a reflective access to it. That is, for *e* to serve as a subject's evidence for believing a particular claim, the subject must be in a position to know through reflection alone that *e* is part of her evidence for believing that claim. What I mean here by "reflective access" or "knowledge through reflection alone" is that the subject does not need to carry out any further empirical investigation to know whether she possesses *e* as evidence for believing a particular claim. As states such as perceptual or memorial experiences are typically reflectively accessible to a subject in that sense—a subject can know through introspection alone whether something visually appears to her in a certain way—the view that a subject's evidence is limited to facts about certain

non-factive mental states she is in is particularly well suited to meet this demand on evidence.

This is not to say, however, that one is forced into such a view upon accepting this demand. Epistemological disjunctivists such as McDowell (1995, 1998), Neta (2008), and Pritchard (2011a, 2011b, 2012) who do not restrict evidence in such a way also take evidence to be, at least in the epistemically good cases, reflectively accessible to a subject. Pritchard, for instance, remarks that:

If the facts in virtue of which one's beliefs enjoy a good epistemic standing are not reflectively available to one, then in what sense is one even able to take epistemic responsibility for that epistemic standing?

(2012, p. 2)

In his view, whenever one sees that p , one is in possession of *factive* evidence for believing that p and one can take responsibility for one's knowledge that p precisely because that evidence is reflectively available to one.³ According to epistemological disjunctivists, a conception of evidence that does not limit one's evidence to internal facts is thus not, in principle, incompatible with the claim that a subject enjoys a privileged access to the evidence she possesses.

Nevertheless, such a compatibilist position faces serious challenges. In particular, it seems to generate a paradox similar to the one outlined by McKinsey (1991) relative to content externalism. To see why, let e stand for some factive empirical evidence possessed by a subject for believing that p and assume for now with Pritchard that the subject's factive evidence consists of her seeing that p .⁴ If the subject is in a position to know through reflection alone that she has the factive evidence e for believing that p , then she should be in a position to know through reflection alone that p is true. For she is also in a position to know *a priori* that her having that factive evidence entails that p is true. Yet, this is unacceptable as a subject is not in a position to know empirical truths in such a way.⁵

For Pritchard (2012) the paradox is, however, only apparent. This is due to the fact that Epistemological Disjunctivism is not committed to the view that a subject's factive evidence for believing that p is reflectively accessible in any sort of circumstances. More precisely, he (2012, pp. 49–52) takes Epistemological Disjunctivism to be only committed to the view that a subject's factive evidence for believing that p is reflectively accessible to her in epistemically favorable circumstances. That is, in circumstances in which the subject believes that p on the basis of her factive empirical support. If, for instance, a subject sees that someone is crossing the field where she is standing and, yet, does not believe that someone is crossing the field because of some misleading defeaters in her possession, Pritchard sees no reason to concede that she has a privileged access to her factive empirical support. For:

we shouldn't judge the good cases—i.e., the good+ cases [cases where the subject believes that p on the basis of her factive empirical support]—by the lights of the not-so-good cases, and the merely good case [cases where the subject fails to believe that p on the basis of her factive empirical support], while not quite a bad case, is certainly a not-so-good case.

(2012, p. 51)

Now, if cases where a subject has a reflective access to her factive empirical evidence for believing that p are limited to cases where she actually believes that p on the basis of that evidence, the problematic conclusion just considered does not seem to follow from the premises. In epistemically favorable circumstances, the subject knows that p on the basis of her factive empirical evidence for p and not on the basis of reflection alone. What she is in a position to know through reflection alone is that she is in possession of factive empirical evidence for believing that p .⁶ Hence, in Pritchard's view, it simply does not follow from the claim that one can have a reflective access to one's factive empirical support that one can know, through reflection alone, empirical claims.

But even if one grants Pritchard that Epistemological Disjunctivism does not have as a consequence that one can know empirical truths through reflection alone, the compatibilist position still faces the challenge, outlined by Kelp and Ghijsen (2016) as well as by Ranalli (2019), of explaining how a subject can enjoy a reflective access to her factive empirical evidence. Suppose that a subject sees that p and that, as Pritchard thinks, her seeing that p consists of her factive empirical evidence for believing that p or, as I claimed in the previous chapter, that p itself constitutes her evidence. What could put that subject, provided that she is in epistemically favorable circumstances, in a position to know through reflection alone that her seeing that p or p itself is part of her evidence for believing that p ? As Pritchard acknowledges, cases in which a subject sees that p and cases where it merely appears to her that p are not introspectively distinguishable. A subject cannot know through introspection alone whether she sees that someone is crossing the field where she is standing.⁷ Consequently, what puts a subject in a position to know through reflection alone that her seeing that p or p itself is part of her evidence for believing that p cannot consist of purely introspectable grounds.

Yet, if what is introspectively available to a subject when she sees that p is not sufficient for her to be in a position to know that her seeing that p or p itself is part of her evidence for believing that p , what is? One option, highlighted by Ranalli (2019), is to rely on the distinction between discriminating support and favoring support drawn by Pritchard (2012). Recall that, according to him, even if a subject's experience does not result from the exercise of an ability to discriminate between one alternative over the other—say the zebra alternative over the cleverly disguised mule alternative—it can

nevertheless provide her with evidence that favors one of the alternatives. If this is correct, as I believe it is, then one could maintain that a subject's knowledge that her seeing that p or that p itself is part of her evidence is not grounded in introspective discriminating support but is instead grounded in non-empirical favoring support. The problem is, as Ranalli (2019, p. 202) notes, that in the case of a subject's seeing that p , it is not at all clear what such non-empirical favoring support could consist of. What could favor, if not one's exercise of certain discriminative abilities, the claim that one sees that p over the alternative that it merely appears to one that p ? Arguing that the fact that the subject sees that p is itself manifest to her in epistemically favorable circumstances would not be of much help, for that would simply move the problem a step further.⁸

The difficulties raised for the compatibilist position defended, *inter alia*, by Pritchard are, however, not, in my view, reasons to abandon this position altogether. They provide, instead, an incentive for rethinking the way in which the reflective access a subject has to her factive evidence should be construed. In particular, I submit that, in epistemically favorable circumstances where a subject sees that p , a subject has a reflective access to her factive empirical evidence for believing that p —*i.e.*, p itself—in the sense that she is in a position to *understand*, solely by reflecting on her situation, why p should be believed.

As outlined by Pritchard, in circumstances where a subject sees that p and believes that p on the basis of her factive empirical support, she has an empirical access to the fact that p . She knows that p on the basis of her factive empirical evidence for p . In addition, the subject has an introspective access to her visual experience as of p . She is in a position to know, based on the exercise of her introspective abilities, that it visually appears to her that p . As a result, in such circumstances, the subject has a knowledge of an evidential situation, which comprises the fact that p and the fact that it visually appears to her that p , that is partly reflective and partly empirical. Now, as noted in the previous chapter, that situation is the normal one in the sense that it does not require any special explanation. And what this means, regarding the problem at hand, is that the subject who finds herself in that situation and has knowledge of it does not need any particular explanation to understand the situation she is in. From her particular perspective—that is, as a subject responsible for the normative standing of her beliefs—it is sufficient for her to reflect on that situation to understand that the reason why the proposition p should be believed is that p is the case.

Note that the type of understanding in terms of which I propose to construe the reflective access a subject has to her evidence in favorable epistemic circumstances is importantly different from the one discussed in the previous chapters. In such circumstances, the subject does not need an explanation of the situation which she has knowledge of to understand that situation from her particular perspective.⁹ Her evidential situation does not need any particular explanation to be intelligible to her. It suffices that she reflects on

that situation for it to become intelligible to her from her particular perspective. The understanding she is in a position to gain is therefore purely reflective and does not depend on the empirical grounds she has for endorsing a particular explanation of her situation. It is by virtue of the very nature of her situation which is characterized by the fact that p is perceptually *manifest* to her that she is able to understand why it is right for her to believe that p solely by reflecting on that situation.¹⁰

Of course, it is crucial that the subject actually believes that p on the basis of her factive empirical support for her to be in a position to gain such a reflective understanding of her situation. Unless she has access to the situation she is in, which involves knowing that p based on the factive empirical evidence she possesses, she cannot come to understand why p should be believed solely by reflecting on that situation.¹¹ This is because what makes it possible for the subject to reflect on her situation in such a way as to understand why p should be believed is the partly empirical knowledge she has of that very situation. Accordingly, the mere fact that p is perceptually manifest to a subject is not sufficient for her to have a reflective access to her factive empirical support for p . It is the exploitation, by the subject, of that favorable evidential situation that makes it possible for her to gain a reflective understanding of the situation she is in. A subject who sees that p but, because of misleading defeaters, refrains from believing that p is in no position to understand, solely by reflecting on her situation, why p should be believed.

If the proposed construal of the reflective access a subject has to her evidence in epistemically favorable circumstances is correct, then the difficulties just raised for the position defended by Pritchard are due to the fact that that access has traditionally been conceived of in terms of the knowledge one has of one's evidence. Framing the question in terms of the reflective knowledge a subject has to the evidence she possesses immediately raises the question of the reflective grounds for such knowledge. In contrast, framing that question in terms of the reflective understanding a subject can gain of her evidential situation from her particular perspective does not raise the same difficulties; at least not in the same way. For, according to the proposed construal of the reflective access a subject has to her evidence, what is specific about the understanding a subject can gain of her situation when she has access to it is that it does not require any ground for accepting a particular explanation of that situation. Whenever p is manifest to a subject and that subject has secured a cognitive access to her situation, she does not need any particular explanation to understand that situation. By virtue of the very nature of the situation she is in, the subject is in a position to understand why p should be believed without the need of explaining it.

6.3. *New Demons and Mere Appearances*

Let me turn to epistemically sub-optimal circumstances and, in particular, cases where a subject is not in a position to know that p because it *merely*

appears to her that p . In such cases, is the subject even justified in believing that p and, if she is, what kind of access does she have to the evidence she possesses for p ? Consider the New Evil Demon hypothesis: S_i who inhabits our world undergoes experiences that are introspectively indistinguishable from the experiences undergone by S_j who inhabits a world ruled by an evil demon. Things appear to S_i as they appear to S_j and, as a result, S_i holds exactly the same beliefs as S_j . Yet, because S_j is systematically deceived by an evil demon, she, contrary to S_i , does not see that such and such is the case and does not remember that such and such was the case.

Several philosophers share the intuition that S_i and S_j , nevertheless, have the same evidence for the propositions they believe and that, as a result, both subjects are equally justified in believing those propositions.¹² If, as some explanationists have it, the evidence each of the two subjects has is constituted of some of the non-factive mental states they are in, then it naturally follows that both S_i and S_j are equally justified in believing those propositions. After all, there is no introspectable difference between the experiences S_i undergoes and the experiences undergone by S_j . However, according to the considerations put forward in the previous sections, this cannot be the case. The evidence S_i possesses must differ from the evidence possessed by S_j .

This does not entail, however, that S_j lacks evidence altogether for the beliefs she holds. As already noted, I see no reason to deny that S_j 's experiences can contribute toward the justification she has for the beliefs she holds. The propositions believed by S_j , while being false, are part of the best explanation S_j has for the way things appear to her and, therefore, she can be justified, given the experiences she undergoes, in believing those falsehoods. But what kind of access, if any, does S_j have to the evidence she possesses for believing the propositions she endorses? Is that access similar to the one S_i has to the factive evidence she possesses for her beliefs concerning the external world? Plausibly, the facts that justify S_j in believing certain propositions about the world can become manifest to her through introspection. S_j can know, by attending to her experiences, that such and such appears to her to be in a certain way. Consequently, it seems that S_j enjoys a purely reflective access to the evidence she possesses. Yet, given the proposed construal of the reflective access a subject has to the evidence she possesses, things are slightly more complicated.

I argued that in favorable epistemic circumstances where a subject sees that p and believes that p on the basis of her factive empirical support, she enjoys a reflective access to her evidence in the sense that she can, by virtue of the nature of the situation she finds herself in, understand why p is to be believed solely by reflecting on her situation; that situation comprising the fact that p and the fact that it appears to her that p . Now, in unfavorable epistemic circumstances where it merely appears to the subject that p , things are quite different. First, the subject's epistemic access to her situation is only partial, for her situation comprises the fact that $\neg p$ and the fact that it appears to her that p . Having a complete access to that situation would involve, for the

subject, knowing that $\neg p$. Second, the situation the subject is in is abnormal or exceptional, which means that to be understood by the subject from her perspective, that situation requires an explanation. Irrespective of the subject's access to the situation she is in, in such circumstances, she cannot possibly understand that situation solely by reflecting on it. Understanding that situation from her particular perspective requires explaining why it appears to her that p despite p not being the case by means of the citation of interfering factors such as the intervention of an evil demon.

It follows that when it merely appears to a subject that p , her access to her evidence for believing that p is somewhat opaque and resembles the type of access an outsider could have to that evidence. While the subject can know, through reflection alone, that such and such appears to her to be a certain way, further empirical investigation is needed to understand the evidential situation she is in and thereby why it is right for her to believe that p . Again, this does not mean that the subject, in such a situation, lacks evidence altogether for believing that p . It only means that for her to understand the evidential situation she is in, she needs independent empirical grounds to endorse an explanation of that situation.

The opacity of a subject's access to her own evidential situation in cases where it merely appears to her that p accounts, I believe, for the fact the justification she has in such circumstances is akin to an excuse for believing a falsehood. According to philosophers such as Pritchard (2012), Littlejohn (forthcoming), and Williamson (forthcoming), the New Evil Demon hypothesis supports drawing a distinction between justification and epistemic excuse or blamelessness. While S_1 can be regarded, given her evidence, as being justified in believing the propositions whose truth she endorses, S_1 , according to that distinction, is rather excusable or blameless for the beliefs she holds. As I do not see any reason for denying that S_1 's experiences can support the propositions she believes to be true, I will not rely on this strategy to account for the intuitions elicited by the New Evil Demon hypothesis.¹³ Nevertheless, the notion of epistemic excuse is closely connected to and sheds light on the opacity of the access S_1 enjoys to her evidence. The reason why S_1 's justification for believing propositions about the world looks like an excuse is that if S_1 were to understand why it is right for her to believe those propositions, she would thereby understand why believing those falsehoods is excusable for her. To see why, consider a subject to whom it merely appears that p . Understanding, for that subject, why p should be believed (by her) involves understanding that the reason why she should believe that proposition is that it merely appears to her that p . This is because $\neg p$ is central to the understanding of that subject's evidential situation. In such circumstances the subject can understand her evidential situation only by finding out that $\neg p$. This, in turn, explains why, for a subject to whom it merely appears that p , gaining a genuine understanding of why it is right to believe that p amounts to gaining an understanding of why it is excusable, for her, to believe that p . The access the subject has to her own

evidential situation is such that she cannot, at the same time, believe what she has justification for believing and genuinely understand why it is right for her to believe those propositions.¹⁴

In unfavorable epistemic circumstances, a subject thus lacks a proper grip on the reason why p should be believed (by her) even if she is in a position to know, through reflection alone, the facts that constitute her evidence for believing that p . This raises, of course, the question as to whether the subject can really be regarded as being responsible for the normative standing of her beliefs in such circumstances. After all, the demand that what constitutes a subject's evidence be reflectively accessible to her is motivated by the claim that justification and responsibility for the normative standing of one's beliefs go hand in hand.

To address this particular question, let me first insist on the fact that if it merely appears to a subject that p , an explanation of the evidential situation she is in is required for her to understand that situation from her particular perspective. Now, if the subject is justified in believing that p given her overall evidence in such circumstances, then she can expect the explanation stating that p to represent her best attempt at accounting for why it appears to her that p by virtue of containing the true proposition p . As her situation comprises the fact that $\neg p$ and that it appears to her that p , this means that the explanation stating that p represents her best attempt at accounting for the situation she is in. By believing that p on the basis of the evidence she has for p , the subject can therefore be regarded as being responsible for the normative standing of her belief. While it is true that she fails to gain a genuine understanding of her evidential situation what she comes to believe represents her best attempt at understanding, from her particular perspective, why it is right for her to believe that p and this, I suggest, is sufficient to account for her responsibility in the normative standing of her belief.

Coming back to the New Evil Demon hypothesis, it should be clear that S_1 's epistemic situation differs substantially from that of S_2 . While both subjects can be regarded as having evidence for the beliefs they hold, only S_1 's evidence suffices to explain why she is correct concerning the propositions she believes about the external world. In addition, while S_1 is in a position to understand, solely by reflecting on her situation, why the propositions whose truth she endorses are to be believed, S_2 's understanding of the evidential situation she is in is only apparent. For the explanation representing S_1 's best attempt at understanding that situation cannot, as a matter of fact, provide her with such an understanding. What fundamentally distinguishes those two subjects with respect to the access they have to their evidential situations is thus that while S_1 's evidential situation is directly intelligible to her, S_2 's understanding of the evidential situation she is in, due to the abnormality of that situation, has to be mediated by an explanation. An explanation which, to provide a genuine understanding of S_2 's situation, needs to state that what S_2 believes is in fact false.

6.4. *Evidence as What Is Manifest*

The proposed construal of the reflective access a subject has to her evidence manages to show, I believe, in what sense internal and external facts that are manifest to a subject can be reflectively accessible to her. But does a subject's evidence only consist of facts that are manifest to her? Providing a positive answer to that question would involve making two distinct claims; one pertaining to the constituents of a subject's evidence, the other pertaining to the conditions under which evidence is possessed by a particular subject. According to the former claim, a subject's evidence is constituted exclusively of facts or true propositions. According to the latter claim, facts or true propositions are possessed, as evidence, by a subject whenever they are manifest to her.

Why think that only facts or true propositions can constitute the evidence possessed by a subject? According to what I argued in the first two chapters of this book, evidence can be conceived of as what is in need of an explanation for a particular subject. Now of course, there is a sense in which a false proposition can be part of what is in need of an explanation for a subject. If a subject falsely believes that her house is burning, she can be regarded as taking, either tacitly or explicitly, that state of affairs as being in need of an explanation. She is disposed, upon reflection on the content of her belief, to consider potential explanations of why her house is burning and to endorse one of those explanations. In that sense, the false proposition 'my house is burning' is part of what is in need of an explanation for her.

Yet, in the present context, the statement "what is in need of an explanation for S" should not be understood as being equivalent to the statement "what S takes to be in need of an explanation." For only facts can, strictly speaking, be in need of an explanation and it is difficult to see how a false proposition that is taken to be true by a subject could impose substantial constraints on what she should believe by virtue of being potentially explained by some explanation available to her. If a subject's evidence includes the false proposition 'my house is burning' and p is part of the best explanation she has for that evidence, how could that false proposition support believing that p given the explanationist conception of evidential support defended in the first chapters of this book? After all, that the subject's house is burning does not require any explanation. There is no fact to be explained. The reason why the evidence possessed by a subject can be the source of substantial rational constraints on her beliefs, as Neta (2018, pp. 42–43) puts it, is that its constituents are truly in need of an explanation. By forming beliefs regarding the content of explanations that could, if true, explain the constituents of her evidence, a subject adjusts, as Williamson (2000, p. 202) stresses, her beliefs to the truth and it is precisely because the truth requires, in an objective sense, to be explained that it makes sense to regard that subject as being required to believe in accordance to the evidence she possesses. When a subject falsely believes that her house is burning, she

can be regarded, at best, as wrongly taking herself to be in possession of a certain piece of evidence.

The fact that a subject's evidence is restricted to truths is thus plausible given what has been argued regarding doxastic justification and given the nature of the constraints evidence imposes on a subject's beliefs. But what reason is there for thinking that truths can be possessed as evidence only insofar as they are manifest to a subject? Facts that are manifest to a subject can be thought of as facts that are available to that subject to arrive at certain conclusions without the need of inference. They are, in that sense, directly available to the subject and, as already noted, it is plausible that this is an important characteristic of evidence: it consists of what a subject has, in the most basic sense, to go on in arriving at certain views. Suppose, for instance, that a subject sees that p , forms the belief that p and correctly deduces that q from p . Once the subject knows that q , if she were to draw another conclusion r from q it would be natural to regard her evidence for believing that r as consisting of p —*i.e.*, what she saw. For, it seems that, although she knows that q , her being right concerning the fact that r is ultimately explained by the fact that she saw that p . Philosophers such as Bird (2004, 2018), however, offer reasons to resist this intuition which, in their view, leads to unacceptable consequences. Consider the following case slightly adapted from Bird (2004, p. 255):

Soluble Liquid: Suppose that it is important for a subject to know whether the substance in a particular jar is soluble and that she comes to know that it is by inferring it from the fact that it dissolved, as she saw. Later, the subject, while still remembering that the substance is soluble, forgets how she came to know that. For all she can recall, it might have been thanks to someone's testimony. At this point, she comes to learn that the substance in the jar is glucose and so infers that glucose is soluble. What evidence does the subject have for her belief that glucose is soluble?

In Bird's view, if one denies that propositions that are correctly inferred from other known propositions can count as evidence, one is forced to endorse the implausible conclusion that the subject lacks evidence for her belief that glucose is soluble. For the evidence on the basis of which she came to know that the substance in the jar is soluble has been forgotten and, by stipulation, the propositions that could count as her evidence for the belief that glucose is soluble are the propositions 'the substance in the jar is soluble' and 'the substance in the jar is glucose.'

Yet, as noted by Dunn (2014, p. 209), other options are available to the proponents of the view that truths are possessed as evidence only insofar as they are manifest to a subject. In particular, in the case under examination, the subject *remembers* that the substance in the jar is soluble and this seems sufficient for that fact to count as evidence. While the subject's inference allowed her to acquire the knowledge that the substance in the jar is soluble,

the reason why she possesses that proposition as evidence for her belief that glucose is soluble might simply be that that fact is manifest to her by virtue of being remembered. Of course, coming to know, thanks to an inference, that the substance in the jar is soluble is what puts the subject in a position to remember that this is the case in the first place. But I see no reason, given the nature of the case, to explain why she possesses that true proposition as evidence by the fact that she knows it as the result of a correct inference. That is, I see no reason, in light of the fact that inferential knowledge can survive the loss of evidence, to accept Bird's claim that "it is not merely knowledge that gets transmitted by inference but also the status of being evidence" (2004, p. 257).

The conclusion that the subject lacks evidence for her belief that glucose is soluble thus simply does not follow from the view that evidence is limited to facts or true propositions that are manifest to a subject. While paradigmatic examples of facts that are manifest to a subject are facts that are known non-inferentially, learning certain truths thanks to a correct inference can put a subject in a position to remember those truths and, thereby, have them as evidence to form further beliefs. Evidence is constituted by truths that are manifest to a subject because, in the most basic sense, it is what is directly available to her to arrive at certain conclusions. But the fact that some truths have been learned on the basis of an inference does not entail that they cannot be available to a subject as evidence. As suggested by cases of forgotten evidence, memory, rather than inference, is able to confer to some truths the status of evidence although those truths have been learned inferentially.

6.5. Justification, Understanding, and Knowledge

Let me now come back, more generally, to the considerations put forward in this book by first highlighting that the view of evidence discussed in the previous section importantly differs from the one defended by philosophers such as Williamson (2000) according to whom evidence consists of the facts that are known by a subject.¹⁵ If the considerations put forward in Chapter 5 are correct, then that a fact is manifest to a subject does not entail that it is known by that subject. A subject can be regarded as seeing that such and such is the case although misleading defeaters prevent her from exploiting an evidential situation that is such that exploiting it would involve being correct concerning what is perceptually manifest to her.

This is not to say, however, that there is no philosophically interesting connection between the notions of evidence and knowledge. Aside from what has just been argued regarding inferential knowledge, evidence comprises facts that can provide a type of explanatory understanding which is characteristic of situations in which a subject knows something to be the case. The facts that constitute a subject's evidence are the facts that can allow one to understand why that subject is correct concerning certain

claims about the world and the notion of knowledge, according to what I argued, fundamentally tracks situations in which such an understanding can be gained.

Thus, even if evidence cannot simply be conceived of as what is known by a subject, there is an important relation between evidence and knowledge. In situations in which a subject possesses knowledge of reality, facts that are manifest to her and that, as a result, constitute her evidence are the ones that can allow gaining an explanatory understanding of the subject's cognitive access to reality. Those facts, because of the rational demands they impose on the subject, are the ones that can be used to explain her cognitive access to reality. But, as I argued, they are also the facts that, from the subject's own perspective as a believer, delimit what is in need of an explanation.

As outlined in the previous section, to say that the facts that are manifest to a subject delimit what is in need of an explanation for her does not amount to saying that they are what the subject takes to be in need of an explanation. Some unclarity might however persist regarding the precise meaning of this claim. Does it state, for instance, that the subject should try to explain the facts that constitute her evidence as opposed to other facts? This would be, after all, a possible reading that is not equivalent to the claim that a subject's evidence consists of what she takes to be in need of an explanation. Yet, it is far from clear why this claim, understood in such a way, should be accepted. Why should a subject try to explain the facts that are manifest to her as opposed to other facts? Any fact, whether manifest to the subject or not, might require an explanation. One may argue that there is a sense in which the subject should try to understand the facts that are manifest to her as opposed to other facts because one cannot genuinely understand facts that one is not in a position to know. But given what I argued in Chapter 4, there is no reason for thinking that a genuine understanding of why p can be secured by a subject only if she is in a position to know that p . Factual knowledge is not necessary for explanatory understanding and the constituents of a subject's evidence are not special in the sense that she should try to explain those facts as opposed to other facts.

The claim that the constituents of the subject's evidence delimit what is in need of an explanation for her is best understood relative to the justification a subject can have for believing certain claims about the world. The facts that are manifest to a subject are in need of an explanation for her in the sense that, whenever the subject has an explanation for them, those facts alone justify, if certain conditions are met, believing the propositions that are part of that explanation. Saying that the evidence possessed by a subject is what is in need of an explanation for her thus amounts to outlining that facts that are manifest to that subject can support believing propositions whose content goes beyond what is manifest to her insofar as those propositions are part of what could provide an explanation of her evidence. It amounts to outlining that the nature of the rational demands imposed by the truths that are manifest to a subject is tied to the ability of certain

representations of reality available to the subject to contribute to the explanation of those truths and thereby to their understanding.

Doxastic justification is therefore not simply directed at truth. In adjusting her beliefs to the evidence she has, a subject endorses the content of explanations that can be expected to be the right way to account for the facts that constitute that evidence. Note, however, that, as I outlined in Chapter 3, gaining such an understanding depends crucially on the grasp the subject has of the explanations of her evidence that are available to her. Accordingly, it is possible for a subject to adjust, in a perfectly rational manner, her beliefs to the evidence she possesses without gaining any understanding of the facts that constitute the evidence which her beliefs are adjusted to. If a subject completely lacks a grasp of the explanations that are available to her for her evidence, forming, with justification, beliefs concerning the content of those explanations is not sufficient for her to gain an explanatory understanding of her evidence. This is true even if the beliefs she forms are about propositions that are central to the understanding of those facts.¹⁶ The grasping requirement of explanatory understanding is thus critical to seeing in what sense doxastic justification is not simply directed at truth. When a subject adjusts her beliefs to the evidence with a grasp of the explanations of that evidence available to her, the justification she has for holding those beliefs can be regarded as being directed at gaining an understanding of the facts that constitute her evidence. That is, the justification a subject has for believing claims whose content goes beyond what is manifest to her can be regarded as being directed at gaining an explanatory understanding of the facts that constitute her evidence in relation to the subject's knowledge of how to account for some aspects of reality by means of certain representations.

As explanatory understanding involves an accuracy requirement, it should be clear that to regard doxastic justification as being directed at gaining an explanatory understanding of one's evidence in no way conflicts with the claim that doxastic justification is directed at truth. If doxastic justification is directed at gaining such an understanding, then it is also directed at truth for to understand why such and such is the case involves being correct concerning claims that are central to that understanding. But the considerations put forward in this book point to another sense, more fundamental, in which doxastic justification is directed at truth. As facts that are manifest to a subject themselves constitute the propositional justification she has for believing certain claims, adjusting one's beliefs to the evidence one possesses involves adjusting one's beliefs to the truth in quite a literal sense: to believe what is manifestly the case. In addition, according to what I argued regarding the reflective access a subject has to the facts that constitute her evidence, such a literal adjustment to the truth is accompanied by a type of understanding of the subject's own evidential situation that encapsulates the truth-directedness of justification. What a subject is in a position to understand through reflection when she sees, for instance, that p and believes that p based on her factive empirical support for p , is that p being the case is the reason why it is

right for her to believe that *p*. Now, the reason to regard this sense in which doxastic justification is directed at truth as being more fundamental is that a subject's justification for believing claims that could allow her to gain an understanding of the facts that constitute her evidence plausibly depends on the type of access she can have to her own evidential situation. Unless the subject is in a position to understand, without the need of explaining her evidential situation, why the facts that constitute her evidence should be explained in the first place it is not clear how explanatory considerations could have the rational pull they have. The subject's justification for believing claims whose content goes beyond what is manifest to her depends on her being in a position to understand why it is right for her to believe the truths that constitute her evidence in the first place.

The picture of the connection between justification, understanding, and knowledge put forward in this book thus suggests, without assuming that understanding is merely a species of factual knowledge, that the projects of securing knowledge of reality and of gaining an understanding of it are deeply interconnected. The truths that are manifest to a subject and that constitute her evidence impose rational demands whose nature is connected to the knowledge the subject can acquire of certain aspects of reality and to the understanding she can gain of those aspects of reality. They are what, from the subject's own perspective as a believer, stands in need of an explanation and, as a result, supports believing claims which could contribute to the explanation of the constituents of her evidence. But they are also what can allow explaining a subject's cognitive access to certain aspects of reality in circumstances characterized by the fact that the subject is in a position to gain, through reflection alone, an understanding of why the constituents of her evidence stand in need of being explained. Hence, knowledge and understanding are importantly interconnected epistemic phenomena and in the present book, I hope to have shown that their connection has to do with the nature of the rational demands imposed by the aspects of reality that make themselves manifest to us.

Notes

- 1 See for instance Huemer (2001), Conee and Feldman (2004, 2011), and Turri (2009).
- 2 As it should be clear, I do not mean to limit the external facts that are part of a subject's evidence to what she sees or remembers. Other success verbs can aptly describe the evidence a subject has.
- 3 Another reason put forward by Pritchard (2005) to accept this demand on evidence is that unless it is accepted, knowledge is conceived of as being compatible with what he labels reflective epistemic luck. That is, it is possible for a subject to know that *p* although, given what is accessible to the subject through reflection alone, it is a matter of pure luck that her belief that *p* is true. A possibility that is in tension, according to Pritchard, with our ordinary notion of knowledge.
- 4 By *factive empirical evidence*, I mean a constituent of *S*'s overall evidence that is such that it is sufficient to explain why *S* is correct concerning the empirical proposition *p*.

- 5 According to Kraft (2015) this problem is not limited to Epistemological Disjunctivism but actually concerns any non-skeptical internalist theory of justification.
- 6 See Neta and Pritchard (2007) for an earlier presentation of this line of response on behalf of Epistemological Disjunctivism.
- 7 As noted by Ranalli (2019), this epistemic claim is also endorsed by proponents of disjunctive views of the nature of perception such as Martin (2004, 2006). It is therefore independent from metaphysical issues pertaining to the nature of perception.
- 8 See Neta (2002, pp. 669–670) for a similar suggestion.
- 9 One might argue that as the subject does not need an explanation of the situation she is in, the resulting state cannot plausibly be conceived of in terms of understanding. As a matter of fact, my discussion of explanatory understanding in Chapters 3 and 4 suggests that gaining an understanding of a particular phenomenon depends on the possession of an explanation of that phenomenon. When it comes to the understanding a subject has of her own evidential situation, things are however different and this is due to the fact that, in epistemically favorable circumstances, a subject's evidential situation is characterized by certain facts being manifest to her. Whenever p is manifest to a subject and the subject believes that p on the basis of her factive empirical support, she has a route to the understanding of her evidential situation that does not depend on the possession of an explanation of that situation. For her situation, in such circumstances, does not stand in need of an explanation.
- 10 The intelligibility of the subject's evidential situation in circumstances where p is perceptually manifest to her and she believes that p on the basis of her factive empirical support echoes, to some extent, the considerations regarding the grasping requirement of explanatory understanding put forward by Grimm (2014) and discussed in Chapter 3. Recall that for Grimm, understanding is dependent on a direct insight into the way things stand in the modal space which is not itself mediated by an explanation but which, instead, motivates a subject to endorse propositions depicting dependence relations between the objects of that direct modal insight. Here I do not wish to suggest that the intelligibility of the subject's evidential situation in favorable epistemic circumstances is to be construed as a direct insight into modal reality. However, my point is that, in such circumstances, the subject's understanding of why it is right for her to believe that p does not need to be mediated by a particular explanation of her situation and that, as a result, that understanding can be regarded as being the product of a direct insight into the situation she finds herself in.
- 11 The proposed construal of the reflective access a subject has to her factive empirical evidence thus prompts to the same kind of response to McKinsey's style paradoxes as the one offered by Pritchard. Moreover, as that access is not construed in terms of knowledge, it is difficult to see how a paradox which proceeds from known premises to a known conclusion could even arise.
- 12 For an early statement of this intuition see Cohen (1984).
- 13 See Madison (2014, 2018) for arguments directed against this particular strategy.
- 14 If justification comes in degree, then, to account for the intuition that S_1 's justification must somewhat differ from the justification possessed by S_2 , it might be argued that the justification S_1 has is better or stronger than S_2 's justification. While such a claim would be perfectly compatible with what I have claimed so far, I believe that this intuition can be captured by the considerations just put forward regarding the type of access S_1 has to her own evidential situation.
- 15 See also Hyman (1999, 2006), Bird (2004, 2018), and Littlejohn (2011, 2017).
- 16 The possibility just described here is quite a radical one in that a subject who completely lacks a grasp of the explanations of her evidence that are available to

her would be operating with absolutely no knowledge of how to account for what is manifest to her by means of certain representations of reality. In fact, it is not even clear that such a possibility exists for there are reasons to think that there is a deeper connection between the evidence one can possess and the knowledge one has of how to account for some aspects of reality by means of certain representations. Hanson (1961), for instance, claims that there is an important difference between what the layman can see and what the physicist can see. While a layman and a physicist can be visually aware of the same thing when they are, say, looking at an X-ray tube overheating, they are not aware of it in the same way. The layman, contrary to physicist, does not see an X-ray tube overheating. She cannot make, as Hanson puts it, the same use as the physicist of what she is visually aware of. If this is correct, then it is reasonable to think that the knowledge one has of how to account for some aspects of reality by means of certain representations determines, at least partly, what facts are manifest to one.

References

- Bird, A. (2004). Is Evidence Non-inferential? *Philosophical Quarterly*, 54, 252–265.
- Bird, A. (2018). Evidence and Inference. *Philosophy and Phenomenological Research*, 96, 299–317.
- Cohen, S. (1984). Justification and Truth. *Philosophical Studies*, 46, 279–295.
- Conee, E. & Feldman, R. (2004). *Evidentialism*. New York: Oxford University Press.
- Conee, E. & Feldman, R. (2008). Evidence. In Q. Smith (Ed.). *Epistemology: New Essays*. Oxford: Oxford University Press, 83–104.
- Conee, E. & Feldman, R. (2011). Replies. In T. Dougherty (Ed.). *Evidentialism and Its Discontents*. New York: Oxford University Press, 428–501.
- Dunn, J. (2014). Inferential Evidence. *American Philosophical Quarterly*, 51, 203–213.
- Grimm, S. (2014). Understanding as Knowledge of Causes. In A. Fairweather (Ed.). *Virtue Epistemology Naturalized: Bridges between Virtue Epistemology and Philosophy of Science*. Dordrecht: Springer, 329–345.
- Hanson, R. (1961). *Patterns of Discovery*. Cambridge: Cambridge University Press.
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. New York: Rowman and Littlefield.
- Hyman, J. (1999). How Knowledge Works. *Philosophical Quarterly*, 49, 433–451.
- Hyman, J. (2006). Knowledge and Evidence. *Mind*, 115, 891–916.
- Kelp, C. & Ghijzen, H. (2016). Perceptual Justification: Factive Reasons and Fallible Virtues. In C. Mi, M. Slote & E. Sosa (Eds.). *Moral and Intellectual Virtues in Western and Chinese Philosophy: The Turn toward Virtue*. New York: Routledge, 164–183.
- Kraft, T. (2015). Epistemological Disjunctivism’s Genuine Access Problem. *Theoria*, 81, 311–332.
- Littlejohn, C. (2011). Evidence and Knowledge. *Erkenntnis*, 74, 241–262.
- Littlejohn, C. (forthcoming). A Plea for Epistemic Excuses. In F. Dorsch & J. Dutant (Eds.). *The New Evil Demon*. Oxford: Oxford University Press.
- Littlejohn, C. (2017). How and Why Knowledge is First. In A. Carter, E. Gordon & B. Jarvis (Eds.). *Knowledge First*. Oxford: Oxford University Press, 19–45.
- Madison, B. (2014). Epistemological Disjunctivism and the New Evil Demon. *Acta Analytica*, 29, 61–70.

- Madison, B. (2018). On Justifications and Excuses. *Synthese*, 195, 4551–4562.
- Martin, M. (2004). The Limits of Self-Awareness. *Philosophical Studies*, 120, 37–89.
- Martin, M. (2006). On Being Alienated. In T. Gendler & J. Hawthorne (Eds.). *Perceptual Experience*. Oxford: Oxford University Press, 354–411.
- McCain, K. (2014). *Evidentialism and Epistemic Justification*. New York: Routledge.
- McDowell, J. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research*, 55, 877–893.
- McDowell, J. (1998). *Meaning, Knowledge, and Reality*. Cambridge: Harvard University Press.
- McKinsey, M. (1991). Anti-individualism and Privileged Access. *Analysis*, 51, 9–16.
- Neta, R. (2002). S Knows That P. *Nous*, 36, 663–681.
- Neta, R. (2008). In Defense of Disjunctivism. In A. Haddock & F. Macpherson (Eds.). *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press, 311–329.
- Neta, R. (2018). Your Evidence Is the Set of Facts That Are Manifest to You. In V. Mitova (Ed.). *The Factive Turn in Epistemology*. Cambridge: Cambridge University Press, 32–49.
- Neta, R. & Pritchard, D. (2007). McDowell and the New Evil Genius. *Philosophy and Phenomenological Research*, 74, 381–396.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2011a). Epistemological Disjunctivism and the Basis Problem. *Philosophical Issues*, 21, 434–455.
- Pritchard, D. (2011b). Evidentialism, Internalism, Disjunctivism. In T. Dougherty (Ed.). *Evidentialism and Its Discontents*. Oxford: Oxford University Press, 235–253.
- Pritchard, D. (2012). *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Ranalli, C. (2019). Epistemological Disjunctivism and Introspective Indiscriminability. *Philosophia*, 47, 183–205.
- Turri, J. (2009). The Ontology of Epistemic Reasons. *Nous*, 43, 490–512.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, T. (forthcoming). Justifications, Excuses, and Sceptical Scenarios. In F. Dorsch & J. Dutant (Eds.). *The New Evil Demon*. Oxford: Oxford University Press.

Index

- accuracy requirements 5, 54, 59–60,
120; of understanding 60–62
- Aristotle 1–2, 6n1
- Audi, R. 88, 103n7
- awareness 122–123n16
- backward causal explanations 104n17
- barn façade case 69–70, 87, 104n20
- Bayesian conditionalization 103n8
- beliefs 87; about *p* 78–79, 83–84, 94,
103n10; about the future 17–18;
and evidence 116; evolving 37, 120;
improbable 22–24; and justification
9, 21–22, 107; *vs.* knowledge 71; and
understanding 59, 65–66
- Bird, A. 117–118
- Blech, I. 72
- Boyd, K. 70–72
- brain in a vat hypothesis 99–101
- Bratman, M. 65
- Byerly, R. 17–18, 22–23, 26n19
- Cahn, J. W. 72
- Carmen (Nero and the Firefighters
example) 74–77
- Cassam, Q. 86
- causal explanations 45–47, 56n12, 86
- causalmechanistic explanations 11
- centrality 62–63, 66, 80n5; and
idealized assumptions 63–64; and
understanding 66–67
- certainty, and epistemology 2
- Cevolani, G. 37–38
- Chisholm, R. 69, 77
- Clifford, W. 8–9
- Cohen, J. 64–65
- Cohen, S. 95
- Coliva, A. 99–100
- Conee, E. 4, 9–12, 25n10, 107
- content externalism 109
- context: and epistemic appraisal 95–96;
and understanding 66
- contextual assumptions 65–66, 98
- contextualism 95–96, 98
- Craig, E. 1
- Dellsén, F. 52, 57n15, 65
- dependence relations, and
explanations 51
- DeRose, K. 95
- Descartes, R. 6n1, 99
- discriminating support 110–111
- disposition to believe 88
- dispositional beliefs 88
- doxastic justification 9, 42, 117,
120–121
- Dretske, F. 86, 95–96, 98
- Dunn, J. 117
- Elgin, C. 42–43, 50, 59, 64–65, 67, 88
- entailment 18–19
- environmental luck 69–71, 73, 77, 79
- epistemic access 101–102, 108
- epistemic appraisal 95–96
- epistemic luck 69–70, 121n3
- Epistemological Disjunctivism 109
- epistemology 1; and certainty 2; and
understanding 2
- evidence 6, 9, 83, 93, 118–119,
122–123n16; accessibility of
109; and conclusions 17–21;
experiences as 107; explanationist
approaches to 9–10, 12, 25n10,
26n19; and explanations 85–90;
and J-Explanationism 24; known
propositions as 117–118; New Evil
Demon hypothesis 113–115; and
p 9–10; questions about 9; and

- reasoning 8; reflectively accessible evidence 108–112, 114, 116; “seeing” as 85–90, 93, 96–97; truths as 117
- evidential support 83, 89
- explanationism 37–39, 107; and knowledge 92–95; and skeptical paradoxes 95
- explanations 10–11, 25n8, 29, 42, 66; and dependence relations 51; and evidence 85–90; and facts 119; grasping 43–44, 50–52; and inferential knowledge 45–46, 49–50, 55; and justification 14, 42; as propositions 43–44; *see also* *e* variable
- explanatory considerations 29–30; and probability 30
- explanatory privilege 91–92
- explanatory understanding 5, 55nn1–2, 68–69, 79, 120; Grimm on 48–49, 122n10
- factive evidence 109, 111–113, 121n4, 122n10
- facts 8, 11–12, 16, 85–86, 94, 96–98, 102, 107–108, 116, 118–119
- false beliefs 116–117
- falsehoods *see* idealized assumptions
- favorable evidential situations 111–112
- favoring support 110–111
- Feldman, R. 4, 9–12, 25n10, 107
- Foley, R. 30
- Friedman, M. 51
- Fumerton, R. 11–12
- future beliefs 17–18
- Gemes, K. 32–33
- Gettier, E. 39, 40n12
- Gettier cases 38–39, 40n12, 69–71, 85
- Gettier luck 78–79
- Ghijzen, H. 110
- Goldman, A. 11–12, 14, 69–70, 94, 103n13, 104n16, 104n20
- Gordon, E. 63, 103n5
- grasping 43–44, 48, 55, 76, 103n6; and accuracy 54; explanations 43–44, 50–52; H 47; and justification 52–55, 120; modal reality 49; and “seeing” 88; *see also* understanding
- Gratias, D. 72
- Greco, J. 1, 47, 103n12
- Grimm, S. 48–50, 56n6, 56n12, 79, 122n10
- Hanson, R. 122–123n16
- Hempel, C. 86
- Hills, A. 52–54
- Hume, D. 6n1, 8
- ideal gas law ($PV=nRT$) 63–65, 80n7
- idealized assumptions 65, 80n7; and Centrality 63–64
- improbable beliefs 22–24
- inferential knowledge 5, 45–46, 49–50, 55, 104n18, 118
- intuition 47, 60, 122n14
- Jenkins, C. 85, 93–94
- Jeremy hypothesis 22–24, 26n19, 29, 37–38, 89
- J-Explanationism 4–5, 12–13, 16–17, 19–20, 24, 25n13, 33–34, 36–37, 39, 89–90, 101; and the Jeremy hypothesis 23, 29, 37–38; and the truthlikeness hypothesis 36
- justification 2, 83, 122n14; and belief 9, 12, 21–22, 107; doxastic 9, 42, 117, 120–121; and experiences 107; and explanations 14; and grasping 52–55, 120; and misleading defeaters 104n21; New Evil Demon hypothesis 113–115; and probability 22; propositional 9–10, 12–13, 24n4; and understanding 59
- Kelp, C. 53, 66, 70, 110
- K-Explanationism 84–85, 92–95, 98, 100, 102–103nn2–4, 103n12, 107
- Khalifa, K. 49, 55n2, 70, 73–79
- Kim, J. 50–51
- Kitcher, P. 51
- know-how 56–57n14
- knowing 83; and understanding 3–4, 45, 68–79
- knowledge 6, 69–70, 84; and evidence 122–123n16; and explanationism 92–95; inferential 5, 45–46, 49–50, 55, 104n18, 118; and justified beliefs 85; from reflection 108–109; and understanding 121
- knowledge ascriptions 95
- Kraft, T. 122n5
- Kuipers, T. 40n2
- Kvanvig, J. 3, 42, 60–61, 69, 79
- Lawler, I. 65
- legal standards of proof 20
- Lewis, D. 95

- Lipton, P. 39–40n1, 55n2, 96
 Littlejohn, C. 114
 luck 69–71, 73, 77–79
- Makinson, D. 15–16
 Martin, K. 18, 22–23, 26n19
 McCain, K. 9, 11, 13–14, 18, 25n10, 26n19, 107
 McDowell, J. 11, 86–87, 109
 McKinsey, M. 109, 122n11
 Miller, D. 32
 misleading defeaters 104n21, 112
 Mizrahi, M. 80n7
 modal reality, grasping 49
- near neighborhood strategy 18–19
 Nero and the Firefighters example 69–73, 75–78
 Neta, R. 11, 90–91, 93–95, 109, 116
 New Evil Demon hypothesis 113–115
 Newman, M. 5, 44–46
 Newton's Apple example 94
 Niiniluoto, I. 37, 40n2
 Nolan, D. 94
 nomological arguments 11
 non-factive mental states 107–109
 non-skeptical periods 2, 4
 non-statistical evidence *vs.* statistical evidence 20, 26n18
 normalicy 91–92, 103n9
- Plato 6n1; *Theaetetus* 3
 Popper, K. R. 5, 31–33
 Poston, J. 9, 15, 17
 Preface Paradox 15–17, 21–23, 29, 38
 Pritchard, D. 3, 6, 43, 46, 48, 69–70, 77–78, 87, 96–97, 109–111, 114, 121n3, 122n11
 probability: and evidential support 35–36; and explanatory considerations 30; and J-Explanationism 37–38; and justification 22; and the truthlikeness hypothesis 31
 propositional justification 9–10, 12–13, 24n4
 propositional knowledge 83
 propositional understanding 103n5
 P-Truthlikeness 32–33; *see also* truthlikeness formula; truthlikeness hypothesis
 Putnam, H. 99
- radical skeptical hypotheses 99–101
 Ranalli, C. 87–88, 110–111
- reality: accounting for 2; and acting/thinking 1; and skeptical arguments 3, 101–102; understanding 4
 reasoning, and evidence 8
 reflectively accessible evidence 108–112, 114, 116, 122n11
 “remembering” 85–86, 93, 117–118
 Rieber, S. 93–96, 104n15, 104n20
- Schurz, G. 5, 32–33, 38, 40n7, 40n12
 “seeing” 85–90, 93, 96–97, 102; *see also* visual appearances
 Shechtman, D. 72
 skeptical paradoxes 100, 104n24; and explanationism 95; *see also* Zebra example
 skeptical periods 2; *see also* certainty
 skepticism 99
 skeptics 2–3, 99
 Smith, M. 91–92, 103n9
 Sosa, E. 103n12
 Stanley, J. 56–57n14
 statistical evidence *vs.* non-statistical evidence 20, 26n18
 Strevens, M. 10, 44
 Stroud, B. 86
- subjects (S) 1–6; beliefs 78–79, 83–84, 96–98, 120, 122n9; and evidence 8, 93, 119; and facts 119–120; guiding 56n11; and justification 83, 89, 120–121, 122n14; knowing *p* 93; and knowledge attribution 83–84; New Evil Demon hypothesis 113–115; and radical skeptical hypotheses 99; seeing 90–91; and understanding 65
 success verbs 85–86, 93, 102; *see also* “remembering”; “seeing”
 SW-Truthlikeness 33–34, 40n10; *see also* truthlikeness formula; truthlikeness hypothesis
- theory of evolution 10
 Tichý, P. 32
 Tichý-Miller paradox 32
 true propositions (*p*) 1–3, 40n8, 55–56n3, 80nn7–8, 116; beliefs about 78–79, 83–84, 94, 103n10, 120–121; believing 16–17, 66; and e_H 60–62; and the *e* variable 13–14, 33–36; and evidence 9–10, 12–13; evidential support for 35–36; and explanations 11–14; and J-Explanationism 37; and K-Explanationism 93–94; and normalicy 91–92; and remembering

- 92–93; and “seeing” 86, 89–92;
and truthlikeness 32–33, 36–37;
understanding 3, 111–112
truthlikeness formula 31
truthlikeness hypothesis 5, 30–31, 36,
39, 40n3; *see also* J-Explanationism
Turri, J. 86–87
- understanding 5, 80n3, 103n5, 111–112,
122n9; and accuracy 60–62; and
beliefs 59; and centrality 66–67; and
context 66; and environmental luck
71, 73, 77, 79; and epistemology 2;
and evidence 83; and grasping 44,
55, 76; and how/why questions 46,
53; and idealized assumptions 65;
and justification 59; and knowing
3–4, 45, 68–79; and knowledge 121;
lack of 43; and reflective access to
knowledge 111; *see also* explanatory
understanding; Nero and the
Firefighters example
- van Fraassen, B. 95
visual appearances 91; *see also*
“seeing”
- Weingartner, P. 5, 32–33, 40n7, 40n12
Wilkenfield 66
Williamson, T. 1, 9, 24n1, 56–57n14,
86, 114, 116, 118
Wright, C. 100
- Zagzebski, L. 2–4, 103n12