



THE UNIVERSITY *of* EDINBURGH

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

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

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

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

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

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



THE UNIVERSITY
of EDINBURGH

**THE CONCEPT OF GRADABLE
KNOWLEDGE**

CHANGSHENG LAI

**PhD Philosophy
The University of Edinburgh
2019**

Declaration

I, Changsheng Lai, hereby declare that this thesis has been composed solely by myself and that it has not been submitted, in whole or in part, in any previous application for a degree. Except where states otherwise by reference or acknowledgment, the work presented is entirely my own.

Portions of Chapter Two (section 2.1) have been published in *Philosophia*:

Lai, C. (2019). Is 'Knowing that P' Identical with 'Knowing that "P" Is True'?. *Philosophia*: 1-18.

Signature:

Date: 12th June, 2019

Acknowledgement

First and foremost, my greatest gratitude goes to my two supervisors, Duncan Pritchard and Martin Smith. Duncan has been supervising me since I studied at the University of Edinburgh as a master student. He patiently guided me with invaluable comments and timely feedback throughout the whole process of my conceiving, composing, and revising this thesis. Martin carefully read the draft of my thesis and has provided detailed feedback. I benefitted enormously from numerous discussions between us over the past three years. In addition, I am very grateful to J. Adam Carter, who used to be my secondary supervisor before he joined the University of Glasgow, for providing constructive comments on the plan and the first draft of my thesis. The supportive supervisory team made this thesis possible.

I owe sincere thanks to many people who offered invaluable suggestions for improving this thesis, including but not limited to: James Brown, Fausto Carcassi, Michel Croce, Gang Deng, Sanford Goldberg, Modesto Gómez Alonso, Emma Gordon, Peter Graham, Donghui Han, Rie Iizuka, Matt Jope, Luke Kersten, Tony Lau, Chang Liu, Hadeel Naeem, Jim Pryor, Brian Rabern, Andrés Soria Ruiz, Guido Tana, Timothy Williamson, Feng Yu, Lian Zhou, Joost Ziff, and the audiences at conferences in Beijing, Taipei, Shanghai, and Edinburgh.

Last but not least, I would like to express my deep gratitude to my family members for their generous help and warm support during my three years of PhD programme, especially my mother Xiumei He and my elder sister Jiena Lai. Special thanks to my girlfriend Yijie Shen for her love and accompany during the past five years.

Thanks to China Scholarship Council for funding my PhD studies.

In memory of my father, an eminent photographer, artist, Jianzhong Chen.

Table of Contents

ABSTRACT	6
LAY SUMMARY	8
INTRODUCTION	10
CHAPTER ONE: EPISTEMIC ABSOLUTISM AND EPISTEMIC GRADUALISM	17
1. ABSOLUTISM AND GRADUALISM.....	17
1.1: <i>Taxonomy</i>	17
1.2: <i>Literature on Gradualism</i>	33
2. MINIMAL GRADUALISM AND THE EQUIVOCAL INTUITION.....	36
3. ARGUMENTS FOR EPISTEMIC ABSOLUTISM.....	42
3.1: <i>Argument from Linguistic Evidence</i>	43
<i>Jason Stanley</i>	44
<i>Julien Dutant</i>	45
3.2: <i>Argument from Object</i>	49
3.3: <i>Argument from Contextualism</i>	54
4. CONCLUDING REMARKS.....	55
CHAPTER TWO: WHY ABSOLUTISM IS ILL-GROUNDED	57
1. OBJECTIONS TO THE ARGUMENT FROM LINGUISTIC EVIDENCE.....	57
1.2: <i>Objection from Counterevidence</i>	61
1.3: <i>Objection from Non-English Language</i>	62
1.4: <i>Objection from Methodology</i>	65
2. OBJECTIONS TO THE ARGUMENT FROM OBJECT.....	78
2.1: <i>Objection to the KT-Schema</i>	82
2.2: <i>Objection to Other Premises</i>	90
3. OBJECTIONS TO THE ARGUMENT FROM CONTEXTUALISM	93
4. CONCLUDING REMARKS.....	94
CHAPTER THREE: HETHERINGTON ON EPISTEMIC GRADUALISM	95
1. THE FIRST VERSION.....	95
1.1: <i>Introduction</i>	95
1.2: <i>Discussion</i>	104
2. THE SECOND VERSION	109
2.1: <i>Introduction</i>	109
2.2: <i>Discussion</i>	114
3. CRITICISM OF HETHERINGTON'S FIRST GRADUALIST THEORY	117
3.1: <i>Anti-Justificationism</i>	117
3.2: <i>Threshold Absolutism</i>	126
4. CRITICISM OF HETHERINGTON'S SECOND GRADUALIST THEORY.....	133
5. CONCLUDING REMARKS.....	137
CHAPTER FOUR: REMODEL THE DEBATE, THRESHOLD AND SPECTRUM	139
1. THE CURRENT MODEL OF THE DEBATE	139
2. THRESHOLD VS. SPECTRUM.....	147
2.1: <i>Threshold Concepts</i>	147
2.2: <i>Spectrum Concepts</i>	153
3. REMODELLING THE DEBATE	159
4. TOWARDS A UNIFIED ACCOUNT OF KNOWLEDGE	164

5. CONCLUDING REMARKS	171
CHAPTER FIVE: GRADUALISM AND EXPLANATORY POWER	173
1. GRADUALISM AND EPISTEMIC LUCK	174
1.1: <i>Two intuitions</i>	174
1.2: <i>Goldberg's Reconciliation</i>	177
1.3: <i>Absolutism's Aporia</i>	186
2. GRADUALISM AND FAULTLESS DISAGREEMENT	198
2.1: <i>The Puzzle of Faultless Disagreement</i>	199
2.2: <i>Extant Proposals</i>	201
2.3: <i>The Gradualist Solution</i>	206
3. CONCLUDING REMARKS.....	212
CHAPTER SIX: EPISTEMIC GRADUALISM AND SCEPTICISM	213
1. THE SCEPTICAL PARADOX	214
2. THRESHOLD AND SCEPTICISM	221
2.1: <i>The Argument from Reference</i>	222
2.2: <i>The Argument from Petitio Principia</i>	243
3. DISCUSSIONS.....	249
<i>Objection 1</i>	249
<i>Reply 1</i>	249
<i>Objection 2</i>	250
<i>Reply 2</i>	250
<i>Objection 3</i>	250
<i>Reply 3</i>	250
<i>Objection 4</i>	251
<i>Reply 4</i>	251
4. CONCLUDING REMARKS.....	258
CONCLUSION: TOWARDS THE CONCEPT OF GRADABLE KNOWLEDGE	259
BIBLIOGRAPHY.....	265

Abstract

An orthodox view in epistemology holds that propositional knowledge is an absolute ‘yes or no’ affair, *viz*, propositional knowledge is ungradable. Call this view *epistemic absolutism*. This thesis purports to challenge this absolutist orthodoxy and develop an underexplored position—*epistemic gradualism*, which was initially proposed by Stephen Hetherington. As opposed to epistemic absolutism, epistemic gradualism argues that propositional knowledge can come in degrees.

This thesis will examine motivations for endorsing absolutism and then, drawing on Hetherington’s original objections to absolutism, prove that absolutism is ill-grounded. In particular, I will explain why the primary ground for insisting absolutism, to wit, linguistic evidence from ordinary English language, fails to entail that knowledge—that is an ungradable concept. After that, I will revisit Hetherington’s two versions of gradualist theories—both will be revealed to be defective. Moreover, the current model of the debate between absolutism and gradualism constructed by Hetherington will give rise to an equivocal attitude towards the gradability of knowledge. That is, there is a prevailing equivocal view which agrees that knowledge can be improved by virtue of better justification but denies that knowledge is, by and large, a gradable concept.

This thesis proposes to remodel the debate between absolutism and gradualism by basing it on a dispute about whether knowledge has a cut-off point distinguishing knowledge from everything that falls short of knowledge. Succinctly put, whether propositional knowledge has a *threshold*. It will be argued that gradualism, so interpreted, should deny that knowledge has a threshold, and treat knowledge as a *spectrum concept* analogous to ‘red’, ‘warm’, and so forth.

The theoretical merits of this new model of the debate and the reconstructed gradualism will be shown. With a better-constructed gradualist account of knowledge in play, I will demonstrate how gradualism enjoys advantages over absolutism by illustrating gradualism's potential applications in solving epistemological issues that absolutism finds difficult to address. For example, issues related to epistemic luck, faultless disagreements, scepticism, and the relationship between different types of knowledge.

Lay Summary

Is knowledge of a proposition as gradable as beliefs or justification? In ordinary language, we usually do not use ‘knows’ as a gradable term. This motivates *epistemic absolutism*, a standard view in the contemporary epistemology that propositional knowledge is an absolute ‘yes-or-no’ affair. In other words, propositional knowledge is ungradable. This thesis will reveal that this absolutist view is ill-grounded, and I will defend a relatively underexplored view, namely, epistemic gradualism. As opposed to epistemic absolutism, epistemic gradualism argues that propositional knowledge admits of different degrees.

The first chapter will unpack the notions of epistemic absolutism and epistemic gradualism. It will be clarified that this thesis aims to defend a gradualist view that a proposition can be known better or worse. I will also introduce three main arguments for absolutism, in particular, the one grounded in our ungradable daily usages of ‘knows’. All those arguments will be refuted in the second chapter. The conclusion reached in this chapter will be that absolutism is not as well supported by arguments as the standard view has it.

The third chapter is about the most prominent advocate of gradualism—Stephen Hetherington. His two different gradualist proposals will be critically scrutinised. It will be revealed that Hetherington’s two gradualist proposals are both problematic, and hence a new form of gradualism is required to be constructed. This task will be accomplished in the fourth chapter. I will argue that a better model of the debate between gradualism and absolutism should focus on whether knowledge is a concept that has a threshold, just like ‘empty’, ‘extinct’, or ‘die’. Absolutism answers that ‘no’, while gradualism should argue that ‘yes’. To be more specific, gradualism should interpret knowledge as a spectrum concept analogous to ‘red’, ‘bright’, and ‘cold’.

The benefits of reconstructing the debate in this way will be illustrated in chapter 5 & 6. I will argue that the remodelled gradualism can be applied to address many epistemological issues that absolutism cannot solve. I will show that gradualism can provide us with a unified account of various types of knowledge. In addition, gradualism can satisfactorily explain the compatibility problem of environmental luck and knowledge, and the phenomenon of faultless disagreement. Finally, we can have a better solution to the long-held sceptical problem by appealing to gradualism. In conclusion, gradualism is preferable to absolutism.

Introduction

Epistemic absolutism¹ is an orthodox view in the contemporary epistemology holding that propositional knowledge is an absolute ‘yes-or-no’ affair, *viz*, propositional knowledge is ungradable (see Hetherington 2001). This thesis will reveal that this absolutist orthodoxy is ill-grounded, and I will defend a relatively underexplored view, namely, epistemic gradualism. As opposed to epistemic absolutism, epistemic gradualism argues that propositional knowledge admits of different degrees.

Epistemic gradualism was prominently advocated by Stephen Hetherington, who developed two different gradualist proposals. It will be revealed in this thesis that Hetherington’s two proposals are both problematic, and hence a new form of gradualism is required to be constructed. A new method to argue for the gradability of propositional knowledge will be developed in this thesis. Moreover, this thesis will reconstruct the model of the debate around the gradability of knowledge². It will be argued that the remodelled debate should focus on whether knowledge is a concept that has a threshold, to wit, a cut-off point distinguishing knowledge from everything that falls short of knowledge. Gradualism should argue that knowledge is not a concept that has a threshold. Instead, knowledge is best to be viewed as a spectrum concept (analogous to ‘red’, ‘bright’, and ‘cold’). This relates to what Hetherington calls ‘external gradualism’—a view that he did not choose to defend. Rather, Hetherington (2001) only aimed to advocate what he called ‘internal gradualism’—a view that knowledge can be better or worse in terms of the justificatory strength. According to Hetherington (2001), the central divergence between absolutism and gradualism is whether internal gradualism holds water. This thesis will

¹ This definition of ‘epistemic absolutism’ is inherited from Hetherington (2001). It should be distinguished from the view claiming that there are absolute facts about what belief a particular item of information justifies – which is sometimes also called ‘epistemic absolutism’ and is usually seen as the opposite of the so-called ‘epistemic relativism’ or ‘epistemic non-absolutism’ (see, for example, Boghossian 2006).

² Hereafter, unless otherwise stated, the term ‘knowledge’ in this thesis usually refers to ‘propositional knowledge’.

show that this is a misdiagnosis, which will give rise to an equivocal attitude towards knowledge's gradability. That is, the idea of internal gradualism is not as unorthodox as Hetherington presents it to be. Quite the opposite, many people would be glad to admit that, on the one hand, knowledge is gradable in the sense that justification can be better or worse. Notwithstanding, on the other hand, knowledge is by and large ungradable in some other (perhaps more important) senses (to be elucidated).

In this thesis, I will attempt to reject this equivocal attitude by advocating a remodelled gradualist account of knowledge. I will demonstrate that the remodelled gradualism enjoys significant advantages over absolutism. That is, gradualism can be applied to address many epistemological issues that cannot be satisfactorily solved within the standard absolutist picture. Given that absolutism is also not as well supported as the orthodox view has it, gradualism is ultimately preferable to absolutism, and thus the equivocal attitude can be discarded for good reasons.

This thesis will proceed as follows:

In Chapter One, the concept of epistemic absolutism and epistemic gradualism will be introduced. I will show that the gradability of knowledge can be measured along two orthogonal scales: the quantitative scale and the qualitative scale. Correspondingly, there can be two types of epistemic gradualism: 1) quantitative epistemic gradualism, which holds that a proposition can be known more or less; and 2) qualitative epistemic gradualism, which claims that a proposition can be known better or worse. This thesis only purports to defend qualitative gradualism. In addition, Hetherington's distinction between external and internal gradualism will be elaborated. After this, the prevailing equivocal attitude towards the gradability of knowledge will be spelt out. The second half of Chapter One will introduce three mainstream arguments for epistemic absolutism. The primary argument for absolutism is what I call *the argument from linguistic evidence*, according to

which the fact that we usually do not use 'knows' in a gradable way indicates that knowledge is ungradable. In addition, *the argument from object* (see Hetherington 2005) argues that the object of knowledge-that- p is the truth of p . Now that the truth of p is ungradable, knowledge-that- p is also ungradable. Finally, *the argument from contextualism* claims that gradualism is rooted in contextualism, and hence it should be rejected along with contextualism as contextualism is implausible.

These arguments for absolutism will be debunked in Chapter Two. The first section of this chapter will argue against the linguistic argument for absolutism in three ways. Firstly, linguistic data lending support to gradualism in the English language will be provided and analysed. Secondly, I will demonstrate that the linguistic intuition that absolutists resort to is too English-centred by exhibiting linguistic evidence supporting gradualism in non-English languages, in particular, the Chinese language. Thirdly, it will be argued that the conceptual nature of some epistemic terms (e.g., knowledge) will sometimes conflict with their daily usages in the ordinary language. Therefore, the linguistic argument for absolutism is untenable due to the methodological mistake. The second section of this chapter will block the argument from object for absolutism by rejecting the premise that the object of an item of knowledge-that- p is just the truth of p . Furthermore, it will be argued that one cannot entail the ungradability of an epistemic concept C from the ungradability of C 's object. Finally, the objection from contextualism will be disarmed in section three by showing that epistemic gradualism can be independent of contextualism. The conclusion reached in this chapter is that epistemic absolutism is ultimately an ill-grounded view. As a result, we shall give its opposite, *i.e.*, epistemic gradualism, a more serious consideration.

Chapter Three will critically scrutinise the two versions of gradualism proposed by Hetherington. It will be summarised that there are two notable traits of Hetherington's first gradualist proposal: firstly, Hetherington (2001) advocates a strong anti-justificationism view insisting that justification is unnecessary for

knowledge—mere true belief can suffice to constitute the so-called ‘minimal knowledge’. Secondly, Hetherington insists that there is a clear distinction between knowledge and non-knowledge. The second gradualist proposal of Hetherington weakens his radical anti-justificationist stance but appeals to his theory of how-knowledge, which is inspired by Armstrong’s ‘truthmaker theory’ (see Hetherington 2005, 2011). It is argued that ‘knowing how it is that p ’ is a sufficient and necessary condition of ‘knowing that p ’, and hence how-knowledge that p is equal to knowledge that p . Since how-knowledge that p is gradable, knowledge that p should be seen as gradable as well. The two proposals will both be called into question. It will be argued that the anti-justificationist position of Hetherington’s first proposal is unacceptable as it fails to capture the idea that knowledge is incompatible with mere-epistemic-luck and the intuition that knowledge is distinctively more valuable than true belief. In addition, I will argue that Hetherington’s commitment to external absolutism is not only unnecessary but also harmful for a satisfactory gradualist account of knowledge. After that, a crucial logical lacuna of Hetherington’s argument for his second version of gradualism will be revealed. That is, the equal sign between how-knowledge that p and knowledge that p is drawn too rashly. Now that both of Hetherington’s two versions of gradualism are problematic, we need a new gradualist theory.

Chapter Four will remodel the current debate between gradualism and absolutism. As we have noted before, Hetherington misdiagnoses the central divergence between gradualism and absolutism. In contrast, this chapter will argue that a better way to characterise the debate should be concentrated on whether knowledge is a threshold concept (analogous to ‘die’, ‘pregnant’, and ‘empty’) or a spectrum concept (analogous to ‘red’, ‘bright’, and ‘spicy’). The distinction between a threshold concept and a spectrum concept will be articulated. Also, I will discuss the relation (in particular, differences) between a spectrum concept and a vague concept. The advantages of this reconstructed model will be illustrated. With a new model of the debate in play, the second half of this chapter will start rejecting the equivocal attitude by demonstrating

gradualism's first theoretical virtue. That is, it can provide a better solution to the putative *asymmetry problem*—why is knowledge-that ungradable while knowledge-how and knowledge-wh are widely recognised as gradable? By directly denying the existence of this alleged asymmetry, gradualism can provide us with a unified account of different types of knowledge.

The theoretical merits of gradualism will be further displayed in Chapter Five. There are two conflicting (but both reasonable) intuitions regarding whether knowledge is compatible with the fake-barn-county-style environmental luck. With these two uncompromising intuitions in play, the relevant discussion seemed to reach an impasse. How can we resolve this impasse? I will first examine Sanford Goldberg's proposal that appeals to his new account of epistemic luck which is based on his notion of 'epistemic entitlement' and 'adequate explanation' (see Goldberg 2015). It will be shown that Goldberg's entitlement-based account of luck would not only fail to accommodate the knowledge-ascribing intuition, but also eliminate the concept of environmental luck. After this, I will argue that a gradualist variation of Goldberg's entitlement-based account of epistemic luck can reconcile the two conflicting intuitions in a more satisfactory manner. The second application of gradualism that this chapter will discuss relates to the puzzle of faultless disagreements. It is widely endorsed that there are faultless disagreements on knowledge ascriptions, but the existence of genuine disagreements and the faultlessness of epistemic peers seem to contradict each other. How can we account for the seemingly inconsistent notion of faultless disagreements? This chapter will show that gradualism can offer us a better explanation that avoids defects of extant mainstream proposals such as contextualism and relativism.

In closing, Chapter Six will introduce gradualism's application in solving the sceptical paradox. A diagnosis of the sceptical problem will show that the sceptical paradox is premised on an intuitively appealing claim that one's belief that 'I am not a BIV' does not meet the threshold for knowledge. Trading on our commitment to that claim, sceptics are able to get the sceptical argument

off the ground. An absolutist anti-sceptical approach will attempt to reject this claim by proving that there is a threshold that is met. In contrast, a gradualist approach will straightforwardly deny that knowledge has a threshold. It will be argued that the absolutist approach succumbs to the difficulty of finding a qualified threshold for knowledge and the problem of *petitio principia*. On the contrary, by interpreting knowledge as a spectrum concept, gradualism can block the sceptical argument in a more undercutting way without giving rise to those problems that absolutism is subject to. Therefore, knowledge is best to be viewed as a spectrum concept, which enables a more satisfactory reply to scepticism.

The final conclusion reached in this thesis will be that the reconstructed gradualism is preferable to absolutism in light of those theoretical advantages introduced. Nevertheless, there are still much more issues that can be discussed in regard to the debate between gradualism and absolutism. This debate, I believe, deserves more attention and further exploration.

Chapter One: Epistemic Absolutism and Epistemic Gradualism

Abstract: This chapter will provide a state of the art review of the current debate between epistemic absolutism and epistemic gradualism. The first section will unpack the notion of gradualism and absolutism by elaborating various readings of the two conflicting views. It will be introduced that the gradability of propositional knowledge can be measured along at least two pairs of scales, namely, the quantitative/qualitative scale and the threshold/beyond-threshold scale. I will clarify that the version of gradualism that this thesis aims to defend is a sort of qualitative threshold and beyond-threshold gradualism. The second section will give a brief sketch of three attitudes that people might hold in regard to the absolutism/gradualism debate. It will be revealed that, without being noticed by Hetherington, the majority of epistemologists might be holding an equivocal attitude. The last section will summarise three mainstream approaches to argue for absolutism. They are the argument from linguistic evidence, the argument from the object, and the argument from contextualism.

1. Absolutism and Gradualism

1.1: Taxonomy

There are various categories of knowledge. For example, there is knowledge of persons, *e.g.*:

- (1) Mary knows John.

It is natural to accept that knowledge of persons can have degrees, *e.g.*, it is permissible to say:

(1a) Mary knows John better than Jane does, as John is Mary's best friend.

There is also knowledge of things, for instance:

(2) Jonny knows rock music very well as he is a renowned rock star.

Again, it is commonly admitted that knowledge of things can be better or worse. Nothing is infelicitous if one says that:

(2a) Jonny knows rock music better than his 3-year-old daughter.

Both knowledge of persons and knowledge of things fall into the category of a more general phenomenon, namely, objectual knowledge (see Bengson & Moffett 2011; Chappell 2014; Kvanvig 2009; Brogaard 2016). It is normally granted by epistemologists that objectual knowledge admits of degrees.

Another important category of knowledge is the so-called 'knowledge-how'³, for example:

(3) As one of the best pro-wrestlers in the world, Bryan knows how to wrestle.

³ Some philosophers argue that knowledge-how can also be understood as a peculiar sort of objectual knowledge. For example, Bengson & Moffett (2011) suggest that knowing how to ϕ something means standing in a non-propositional objectual understanding relation to a way of ϕ -ing. Trading on this point, Bengson & Moffett construct their so-called 'nonpropositional intellectualism'.

It is widely endorsed by most of the epistemologists nowadays that knowledge-how can be gradable⁴. Linguistic evidence can be felicitous expressions as follows:

(3a) As one of the best pro-wrestlers in the world, Bryan knows how to wrestle far better than Cena does.

In addition, there are also different types of knowledge-wh, for example:

(4.1) Amy knows where the Italian restaurant is.

(4.2) Tommy knows why they eventually decided to dine at that Italian restaurant.

(4.3) Jimmy knows what the name of the Italian restaurant is.

Similar to those sorts of knowledge introduced above, the gradability of most types of knowledge-wh is commonly recognised as well (see, *e.g.*, Sgaravatti & Zardini 2008)⁵. For instance, the following expressions are grammatical:

(4.1a) Amy knows better where the Italian restaurant is than Jimmy, because Jimmy only knows that the restaurant is located in Vic Street, but Amy knows it is at 101 Vic Street.

(4.2a) Tommy knows why they eventually decided to dine in that Italian restaurant better than Amy does, because Tommy knows more details about how the decision was made.

⁴ Exceptions see Pavese (2017), where Pavese argues that knowledge-how, as well as knowledge-that, does not have degrees.

⁵ It is worth noting that Stanley (2011) rejects the gradability of know-that on one hand, while on the other hand, he argues that know-wh can be reduced to know-that, the former of which is typically granted as gradable. Sgaravatti & Zardini (2008) take this as an inconsistency of Stanley's intellectualist view that know-wh is a kind of know-that.

However, whether the following expression is felicitous or not is more ambiguous:

(4.3)??⁶Jimmy knows what the name of the Italian restaurant is better than Tommy does.

A similar type of knowledge is knowledge-whether, whose gradability is also doubtful:

(4.4) Jeremy knows whether the name of that restaurant is *Locanda De Gusti* or not.

(4.4a)??Jeremy knows better whether the name of that restaurant is *Locanda De Gusti* or not than Tommy does⁷.

In contrast, when it comes to the most-discussed type of knowledge in contemporary epistemology, *viz.* propositional knowledge (aka ‘knowledge-that’), the orthodox view is seemingly unanimous that knowledge-that does not admit of degrees (see Stanley 2005, 2011; Dutant 2007; Ryle 1949; Dretske 1981; Pavese 2017; Sgaravatti & Zardini 2008; Crane 2012; *etc.*). This orthodox view is named *epistemic absolutism* (see Hetherington 2001). For absolutists, a proposition can neither be known less or more, nor better or worse. Knowing a proposition is an absolute yes-or-no affair rather than a

⁶ In this thesis, when reporting sentences as ungrammatical, I will use the amount of question marks to represent to which degree a sentence is infelicitous. That is, sentences starting with two question marks are more linguistically suspicious than sentences starting with one. Sentences preceded by three question marks are more clearly infelicitous than those preceded by two or one.

⁷ Unless otherwise stated, linguistic judgments in this thesis are made largely on the basis of my own linguistic intuition. Admittedly, my intuition and judgments might not be agreed by all of the readers, however, this should not affect my main argument which does not hinge on people’s linguistic intuition. Here, these linguistic data are only used to help to set the stage for introducing the idea of epistemic absolutism and epistemic gradualism—these data do not play any substantial role in my main argument.

matter of degrees. Given absolutism, expressions such as (5.1a) and (5.2a) are untoward:

(5.1) Mike knows that Trump won the 2016 US election.

(5.2) Mike knows that the name of the victor of the 2016 US election is Trump.

(5.1a)???)Mike knows that Trump won the 2016 US election better than James does⁸.

(5.2a)???)Mike knows that the name of the victor of the 2016 US election is Trump better than James does.

The platitude that propositional knowledge is absolute/ungradable seems can be used to explain why the gradability of some other sorts of knowledge is doubtful. For example, the infelicitousness of (5.2a) may partially explain the infelicitousness of (4.3b), as it is seemingly natural to interpret ‘S knows what the name of X is ...’ in terms of ‘S knows that the name of X is ...’

Analogously, as ‘S knows whether *p* or not’ can usually be interpreted as ‘S knows that *p*’ or ‘S knows that not-*p*’ (see Groenendijk & Stokhof 1982; Karttunen 1977; Hintikka 1976; Stanley 2011 ch.2; *etc.*), it seems no wonder that the gradability of knowledge-whether is sometimes as unacceptable as that of knowledge-that—just as shown by (4.4a) and (5.1a).

This traditional view that propositional knowledge cannot come in degrees is what Hetherington refers to as *epistemic absolutism*: ‘Knowledge is absolute, in the sense that it is impossible for a person to have *better*, or to have *worse*, knowledge of a fact’ (Hetherington 2001: 3). The perspective of

⁸ These examples are inspired by similar cases provide by Jason Stanley, such as ‘John knows that Bush is president better than Mary does’ (2005: 40). Stanley reported that himself and every non-philosopher informants that he asked found these sentences acceptable. Note that, here, I put three question marks in front of those utterances to indicate that they are clearly infelicitous in accordance with the absolutist orthodoxy. However, gradualists might argue that the oddness of these utterances can be reduced once we spell out the context in which there are sufficient evidential room for these utterances to be known better or worse (see Hetherington 2011: 212-213).

epistemic absolutism has been taken for granted by many epistemologists for a long period, for example, Sgaravatti & Zardini claim that:

‘It’s a well-known fact that knowledge-wh enjoys some forms of gradability while knowledge-that plausibly does not.’ (Sgaravatti & Zardini 2008: 255)

However, few philosophers have seriously developed systematic arguments for epistemic absolutism, which causes Hetherington to call absolutism a ‘dogmatism’. Perhaps that is because absolutism has been taken to be too uncontroversial to be worth seriously defending. Nevertheless, claims or expressions supporting epistemic absolutism can be broadly found in the literature. For instance, when commenting on Michael Tye’s theory of ‘knowledge by acquaintance’, Tim Crane remarked that:

‘[K]nowing things in the ordinary sense admits of degrees. I know Michael Tye fairly well, but not as well as some other people do. Although I have been to Athens a few times, I do not know it as well as Tye does. This contrasts with propositional knowledge, which (on most conceptions) does not admit of degrees.’ (Crane 2012: 149)

Crane distinguishes knowledge of things and propositional knowledge in terms of their gradability—knowledge of things/persons admits of degrees, while propositional knowledge does not. Similar opinions can be found in Carlotta Pavese’s remarks on Bertrand Russell’s theory of knowledge by acquaintance. According to Pavese (2017), this difference in gradability also constitutes the primary reason why Russell held that knowledge-by-acquaintance cannot be reduced to propositional knowledge. Apart from these, Stanley (2005), Dutant (2007) both argued for epistemic absolutism from the aspect of linguistics.

Apart from these recent remarks, the most well-known *loci classici* as regards epistemic absolutism are from Ryle (1949) and Dretske (1981). When

discussing the difference between knowledge-that and knowledge-how, Ryle argues that:

‘We never speak of a person having partial knowledge of a fact or truth, save in the special sense of his having knowledge of a part of a body of facts or truths. A boy can be said to have partial knowledge of the counties of England, if he knows some of them and does not know others. But he could not be said to have incomplete knowledge of Sussex being an English county. Either he knows this fact or he does not know it. On the other hand, it is proper and normal to speak of a person knowing in part how to do something, i.e. of his having a particular capacity in a limited degree.’ (1949: 59)

Ryle famously argues for anti-intellectualism—a view that knowledge-how cannot be exhaustively reduced to knowledge-that. The quotation above shows that an important reason lending support to Ryle’s anti-intellectualist view is that knowledge-how has degrees while knowledge-that does not. Ryle does not explain much why propositional knowledge cannot have gradability in this quoted paragraph. The most salient reason that he mentions is that propositional knowledge cannot be *partially* known, and thereby a proposition/fact/truth cannot be known *more* or *less*, and therefore knowledge-that cannot be graded. It is not hard to find out that there is something subtly different hidden in Ryle’s understanding of gradability of knowledge-that. Cases introduced before conclude that propositional knowledge is ungradable because it seems infelicitous to say that ‘S₁ knows *better/worse* that *p* than S₂’, while Ryle concludes that propositional knowledge is ungradable because it seems infelicitous to say that ‘S₁ knows *more/less* that *p* than S₂’. This distinction between the *better/worse* sense and the *more/less* sense of gradability suggests that we ought to unpack epistemic absolutism in a more fine-grained manner.

Firstly, we can at least make a distinction between *quantitative gradability* and *qualitative gradability*. Pavese (2017) characterises ‘quantitative gradability’⁹ as follows:

‘The phenomenon of quantitative gradability is the phenomenon whereby ascriptions of the form “S knows how to –” can be sensibly modified by so-called proportional modifiers such as “in part”, “for the most part”, “partly”, “partially”.’ (Pavese 2017: 5)

Pavese then argues, in line with Ryle and numerous other epistemologists, that it is awkward to accept that knowledge-that admits of quantitative gradability. The idea is that it does not make sense to say something like ‘I know in part that Sussex is in England’ or ‘Jack knows partly that today is Friday’. Similarly, Dutant (2007) divides paradigmatic degree modifiers into two categories: *quantitative modifiers* (e.g., very/much/more than) and *qualitative modifiers* (e.g., good/well/better than). The measure of qualitative modifiers, according to Dutant, ‘is a degree of goodness’ (2007: 3). Dutant holds that ‘know-that’ can rarely be modified by quantitative modifiers, however, propositional knowledge attributions can be built with qualitative modifications properly. Moreover, in order to challenge Stanley’s conclusion that ‘know-that’ cannot occur with genuine qualitative comparatives, Dutant also lists many cases (see Dutant 2007: 7 & 25, appendix 8.4) where propositional knowledge attributions are genuinely¹⁰ and felicitously built with qualitative comparatives (e.g., ‘better-than’)¹¹.

⁹ See also Berman (1987), Lahiri (1991; 2000), Beck & Sharvit (2002), and Williams (2000). Bierwisch (1989) also makes a similar distinction between *dimensional* and *evaluative* gradable adjectives.

¹⁰ Stanley concedes that propositional knowledge attributions can sometimes occur with qualitative-comparatives in some *idiomatic constructions* such as ‘... knows that ... better than anyone’. However, because ‘better than anyone’ is just an idiomatic construction, Stanley refuses to admit that comparative construction containing such an idiom can prove the *genuine* gradability of ‘knows’ (see Stanley 2005: 40).

¹¹ However, Dutant still refuses to admit that propositional knowledge can have *qualitative* gradability; his reasons will be introduced later.

By far, we can roughly differentiate at least two kinds of absolutism in terms of the quantitative/qualitative distinction:

[Quantitative Epistemic Absolutism]

Propositional knowledge is absolute, in the sense that a fact or truth cannot be known *in part*; one's knowledge of a given proposition cannot be *augmented*; it is impossible for a person to have *more*, or to have *less*, knowledge of a fact.

[Qualitative Epistemic Absolutism]

Propositional knowledge is absolute, in the sense that a fact or truth cannot be known *with higher quality*; one's state of knowing-that-*p* cannot be *improved*; it is impossible for a person to have *better*, or to have *worse*, knowledge of a fact.

On the contrary, there can be two sorts of gradualism correspondingly:

[Quantitative Epistemic Gradualism]

Propositional knowledge is gradable, in the sense that a fact or truth *can* be known *in part*; one's knowledge of a given proposition *can* be *augmented*; it is *possible* for a person to have *more*, or to have *less*, knowledge of a fact.

[Qualitative Epistemic Gradualism]

Propositional knowledge is gradable, in the sense that a fact or truth *can* be known *with higher quality*; one's state of knowing-that-*p* *can* be *improved*; it is *possible*

for a person to have *better*, or to have *worse*,
knowledge of a fact.

It is worth noting that, the quantitative scale and the qualitative scale are not contradicting each other. One who accepts quantitative epistemic absolutism can also be committed to qualitative epistemic absolutism (likewise, one can endorse both quantitative and qualitative epistemic gradualism). Alternatively, one is permitted to defend quantitative absolutism meanwhile reject qualitative absolutism, and vice versa.

Now, with the two distinctive sorts of absolutism in play, it is not hard to find that the absolutist view presented in Ryle's remark is mainly a sort of *quantitative* epistemic absolutism. In contrast, Hetherington (2001) challenges the absolutist orthodoxy mainly along the *qualitative* scale. He clarifies that he discusses the absoluteness of knowledge-that in a particular sense—that is, 'the sense of its being impossible for a piece of knowledge to be better or worse qua knowledge, to be more or less *clearly* knowledge' (2001: 6). Notice that, the 'more or less' in this quotation does not imply quantitative gradability, but is used to modify the 'clearly' after, *viz*, to modify the degree to which a piece of knowledge can qualify as knowledge. Accordingly, Hetherington urges us to embrace *epistemic gradualism*—a view that knowledge-that is gradable in the sense that we can have better or worse knowledge of a proposition. Hetherington argues that propositional knowledge can be graded in terms of the failability¹² (see Hetherington 2001: 49). The more failable one's knowledge that *p* is, the worse one knows that *p*. Given the close connection between failability and justificatory strength, a succinct corollary of Hetherington's gradualism is that the weaker one's justification for *p* is, the worse one knows that *p*. Following Hetherington, this thesis will focus on the *qualitative* gradability of knowledge-that. I will manage to develop a new

¹² In brief, for Hetherington, 'knowing failably' is a broader concept than 'knowing fallibly' – the latter means that S knows that *p* but S's belief that *p* could have been *false*, while the former means that S knows that *p* but S could have *failed* to do so. Chapter 3 will introduce this distinction in more detail.

qualitative gradualist account of knowledge and leave it open whether propositional knowledge is also quantitatively gradable¹³.

After clarifying this, we can now move forward to Dretske's frequently-cited remarks:

'When talking about people, places and topics (*things* rather than facts), it makes sense to say that one person knows something *better than* another—But factual knowledge, the knowledge *that* something is so, does not admit of such comparisons. If we both know that today is Friday, it makes no sense to say that you know this better than I. In this respect factual knowledge is *absolute*. It is like being pregnant: an all or nothing affair.' (Dretske 1981: 363)

Dretske agrees that there can be better or worse knowledge of a person, a place or a topic, but he denies that knowledge of a fact can be equally gradable. What is the essential difference between knowledge of a topic and knowledge of a fact that leads to such an asymmetry here? A possible explanation is that a fact is just a fact—an absolutely true fact. In light of this, there is no ambiguous gradient ramp between knowledge and not-knowledge—the distinction is sharp and clear. Therefore, to know a fact is to know it absolutely, clearly and perfectly.

In response, Hetherington comments that Dretske confuses the absoluteness of 'the cut-off point/threshold between knowledge and not-knowledge' with the absoluteness of 'knowledge *per se*'. The distinction between these two kinds of absoluteness is made clear by Hetherington as he is committed to the former while rejecting the latter. Hetherington clarifies that:

¹³ It strikes me that *quantitative gradualism* is even more unpopular than qualitative gradualism for the majority of epistemologists. After all, it seems to be hard to conceive how a proposition can be known more or less. Besides, it is far from clear what it means by saying that one *partially* knows a proposition. Given that the extant literature mainly focus on discussions around the qualitative gradability of knowledge—that, the central domain of discourse of this thesis will also be limited to the qualitative scale.

‘In denying that knowledge need be absolute (as I will be doing in this chapter and beyond), I will not be denying that there is a cut-off point—and an absolute one at that—between knowing and not knowing.’
(Hetherington 2001: 6)

Also, Hetherington argues that:

‘We can be absolutists about the difference between knowing and not knowing, without our also having to deny that some knowledge that *p* can be better as knowledge than some other knowledge that *p*.’ (2001: 7)

With these being said, Hetherington admits that he is committed to a form of *local* gradualism that consists of two views:

(i) Knowledge is to be absolutely distinguished from whatever is not knowledge (whatever is external to knowledge). This is because there is an absolute cut-off point between knowing and not knowing,
(ii) But *within* the category of knowing, non-absolutism is true. That is because it is possible that some cases of knowledge that *p* are better as knowledge that *p*, or more clearly knowledge that *p*, than other cases of knowledge that *p*.’ (2001: 7)

Hetherington names view (i) as ‘external absolutism’ and calls view (ii) ‘internal gradualism’. A central notion that defines the internal/external distinction is the ‘cut-off point’ between knowing and not knowing. This cut-off point is also interpreted as the threshold or boundary of knowledge (see Hetherington 2006; Hannon 2017; *etc.*). External absolutism is *external* in the sense that the threshold distinguishing knowledge from everything that is *external to* the category of knowledge is absolute. In contrast, internal gradualism is *internal* in the sense that beyond the threshold for knowledge, to wit, *within* the category of knowledge, an instance of knowledge cannot be better or worse than another instance. Accordingly, we can rephrase Hetherington’s dichotomy and summarise another pair of readings of absolutism:

[Threshold Epistemic Absolutism]

Propositional knowledge is absolute, in the sense that the threshold between knowledge and not-knowledge is absolute and clear.

[Beyond-Threshold Epistemic Absolutism]

Propositional knowledge is absolute, in the sense that even beyond the threshold (no matter being absolute or not) between knowledge and not-knowledge, knowledge cannot be better or worse.¹⁴

Threshold epistemic absolutism reflects an intuition that knowledge is absolute in terms of the absoluteness of its threshold. It is just like when saying 'truth is absolute', we mean that there is a clear cut-off point between truth and falsehood. A proposition is either true or false, there is no penumbral buffer zone between truth and falsehood. Similarly, 'being pregnant', 'defeat', 'jump' are all absolute terms in the same sense. Beyond-threshold epistemic absolutism analyses the absoluteness of knowledge in a different sense, that is, once the standard of knowledge is met, nothing can contribute any more to an instance of knowledge's being knowledge. It is analogous to say one's death is absolute, as once one dies, he/she cannot 'die more'—being dead for a longer time does not contribute to the fact that one is dead. On the contrary, 'jump', albeit being of an absolute threshold, is not as much beyond-threshold-absolute as 'die', because the increase of height of leaping can make one's jump a 'higher jump'.

Correspondingly, there are two sorts of gradualism:

¹⁴ Notice that, this threshold/beyond-threshold distinction assesses the gradability of propositional knowledge along a distinct scale from that of the quantitative/qualitative distinction.

[Threshold Epistemic Gradualism]

Propositional knowledge is gradable, in the sense that the threshold between knowledge and not-knowledge is *not* absolute and clear. Instead, the distinction between knowledge and not-knowledge can be gradient.

[Beyond-Threshold Epistemic Gradualism]

Propositional knowledge is gradable, in the sense that beyond the threshold (no matter being absolute or not) between knowledge and not-knowledge, knowledge can be better or worse.¹⁵

According to Hetherington's interpretation, Dretske is just advocating a version of threshold epistemic absolutism, which is compatible with Hetherington's own beyond-threshold epistemic gradualist view. Moreover, Hetherington does not plan to argue for threshold epistemic gradualism, because he is also a threshold absolutist just like Dretske.

Let me further elaborate on what it means by saying that a threshold is 'absolute' (we will revisit this issue and unpack this notion again in Chapter 4). Admittedly, it seems to be too demanding to locate an absolutely clear cut-off point between any two concepts, and hence threshold epistemic gradualism might seem to be implausible in light of this. Furthermore, Hetherington does argue somewhere else that it is impossible to find out an exact threshold for knowledge that can answer the *boundary problem of knowledge*¹⁶ (see Hetherington 2006). This seems to contradict his commitment to threshold epistemic absolutism. It can even be argued that

¹⁵ Just like quantitative gradualism is compatible with qualitative gradualism (so is quantitative absolutism and qualitative absolutism), threshold epistemic gradualism is compatible with beyond-threshold gradualism. One can defend both threshold and beyond-threshold epistemic gradualism – that is what I will do in this thesis.

¹⁶ The boundary problem is 'the epistemological problem of knowing what is the maximum degree of fallibility that is allowable in knowledge's justificatory component' (Hetherington 2006: 42).

the boundary of any concept is vague, thus so is knowledge. So is threshold gradualism advocating the gradability of knowledge in this trivial manner? No. Rather, I suggest that when rejecting the idea that knowledge has an absolute threshold, threshold epistemic gradualists ought to prove at least two things:

(T1) There is not, and doesn't need to be, a reductive analysis of knowledge that comes in the form of *sufficient and necessary* conditions (*cf.* Williamson's anti-analysis¹⁷ stance).

(T2) There is no absolute and non-arbitrary boundary for how a necessary component of propositional knowledge can be counted as *good enough* to constitute knowledge (*e.g.*, there is no absolute cut-off point between a 'firm-enough' and a 'not-firm-enough' or a 'justified-enough' and a 'not-justified-enough' belief *for constituting knowledge*).

The aforementioned trivial understanding of the non-absoluteness of knowledge's threshold mainly reflects the gist of (T2). Hetherington also argues for the non-absoluteness of knowledge in a similar sense. However, Hetherington's rejection of (T1) makes him a threshold epistemic absolutist. As I will introduce in the third chapter of this thesis, Hetherington endorses that there is an absolute sufficient and necessary condition for knowledge, *viz.* true belief, which constitutes what he refers to as 'minimal knowledge'. The denial of (T1) represents a typical understanding of 'absolutism' in philosophy. For example, *moral absolutism* is generally understood as a view holding that there is a fixed universal sufficient and necessary condition of being morally right (otherwise, morally wrong); *truth absolutism* generally means there is a fixed universal sufficient and necessary condition of being

¹⁷ In this thesis, I will mainly discuss reductive analysis of knowledge.

true (e.g., corresponding to the fact). Therefore, *contra* Hetherington, threshold gradualists should also take (T1) into consideration, as it is supposed to be an essential part of the debate between threshold gradualism and threshold absolutism.

In contrast, beyond-threshold epistemic gradualism seems to be much less controversial than threshold epistemic gradualism, because it is widely granted that one's epistemic support (justification, evidence, reliability, etc.) for one's knowledge-that-*p* can be stronger than someone else's. It seems to be a highly plausible idea that the better one's knowledge-that-*p* is justified (even beyond the threshold), the better one's knowledge is. Hetherington might also be aware of this, so he constructs an alleged dilemma for epistemic absolutists:

- '1. Knowledge is absolute, non-gradational. That is, knowledge that *p* cannot be better or worse as knowledge that *p*.
2. Each part of one's justification for a belief which is knowledge is contributing somehow to the belief's being knowledge, to how it is the knowledge it is.' (Hetherington 2001: 21)

Hetherington argues that, due to the absolutist credo, absolutists have to accept 1, which must be done at the cost of denying 2. But denying 2 seems to be an unpalatable upshot for absolutists, given the intuitive plausibility of 2. Hetherington thus concludes that absolutism is subject to a dilemma.

However, I suspect that absolutists may easily disarm this dilemma by accepting 2 and denying the incompatibility between 1 and 2, and thus embrace a somewhat 'equivocal' stance of the gradability of propositional knowledge. This attitude will be further discussed in the next section. In order to avoid this equivocal attitude, differing from Hetherington's local gradualism (*i.e.*, beyond-threshold gradualism + threshold absolutism), the version of gradualism that I plan to advocate in this thesis is a *global* one. That is to say, I propose to defend both threshold and beyond-threshold epistemic

gradualism and reject both threshold and beyond-threshold epistemic absolutism.

1.2: Literature on Gradualism

Now we have analysed some representative remarks supporting epistemic absolutism in the current literature. Although absolutism is the orthodox understanding of propositional knowledge endorsed by most epistemologists, there are still voices from the minority of epistemologists arguing for gradualism explicitly or implicitly. For example, from the perspectives of contextualism, Unger (1975) argues that 'know' is analogous to *context-insensitive* gradable adjectives¹⁸ that denote a limiting point on the relevant scale (e.g., 'flat' and 'empty'), and Cohen (1999) holds that 'know' is analogous to *context-sensitive* gradable adjectives (e.g., 'tall' and 'large').

Apart from that, Sosa sometimes displays his gradualist tendency on the basis of his theories of virtue epistemology. For instance, his notable cross-kind hierarchical theory of knowledge distinguishing 'animal knowledge', 'reflective knowledge' and 'knowledge full well' can be regarded as an (albeit somewhat peculiar) version of qualitative epistemic gradualism. When discussing the distinction between animal knowledge and reflective knowledge, Sosa wrote that '[k]nowledge seems a matter of degree in a variety of respects' (2001: 194). In his 2007 book, Sosa emphasised again that 'our distinction between two sorts of knowledge, the animal and the reflective. Any full account would need to register how these are matters of degree.'¹⁹ (Sosa 2007: 32) These snippets of quotations, of course, are not sufficient to reflect the position of Sosa's view takes in the gradualism/absolutism debate. We will revisit Sosa's theory and the relation between gradualism and virtue epistemology in Chapter 4.

¹⁸ See also Dutant (2007) for relevant discussions.

¹⁹ Moreover, Sosa (p.c.) pointed out that Stanley's linguistic objection to gradualism failed to notice that knowledge can be graded in terms of quality.

Influential as contextualism and virtue epistemology are, the most renowned and prominent advocate of epistemic gradualism in the current literature, as aforementioned, is still Stephen Hetherington²⁰. In his 2001 book *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*, Hetherington argues for the gradability of propositional knowledge explicitly by rejecting what he calls as ‘two dogmas of epistemology’, *i.e.*, absolutism and justificationism. His interpretation of absolutism has been introduced before. And with regard to another ‘dogma’ he refers to as ‘justificationism’, which represents the doctrine that justification is necessary for knowledge, Hetherington argues that knowledge can be achieved without justification and contends that ‘minimalism’ (mere true belief can sometimes suffice to constitute the minimal knowledge) should be accepted instead. This anti-justificationist view is perhaps even more radical and unorthodox than his gradualist position. While in another paper on gradualism which was published in 2005, Hetherington softens his position and tries to argue for gradualism without insisting his anti-justificationism stance so firmly. Instead, Hetherington (2005) turns to develop gradualism on the basis of his theory of how-knowledge²¹ (*i.e.*, ‘knowledge of how it is that p ’), which is heavily influenced and motivated by the ‘truthmaker theory’ of Armstrong (2004)²². This idea is further developed in his 2011 monograph ‘How to Know’ along with a ‘practicalist account’ of gradualism (we will return to these points in Chapter Three). It is argued that ‘knowing how it is that p ’ is a sufficient and necessary condition of ‘knowing that p ’, so how-knowledge that p is equivalent to knowledge that p . Since how-knowledge that p is gradable, knowledge that p ought to be deemed gradable as well. Both versions of

²⁰ For a helpful summary of other explicit and potential proponents of anti-absolutism, see Hetherington (2011: 51-62).

²¹ For Hetherington, ‘how-knowledge’ is slightly different from ‘knowledge-how’. The former basically means ‘knowing how it is that p ’ rather than ‘knowing how to ...’—although the two concepts are closely related.

²² According the truthmaker theory, knowledge is ordinarily taken to be knowledge of a truth, and every truth has at least one truthmaker, *i.e.*, ‘something in the world that necessitates that truth’s being true’ (Hetherington 2005: 137). Those most discerning truthmakers are ‘minimal truthmakers’ for a truth. For example, minimal truthmakers for ‘Stephen is a male basketball player’ are ‘Stephen’s being male’ and ‘Stephen’s being a basketball player’. According to Hetherington (2005), one’s how-knowledge that p can be graded in terms of how many parts of (a) minimal truthmaker(s) for p are known by the subject.

Hetherington's epistemic gradualism are closely related to his innovative and controversial theory about knowledge-how, which is known as 'radical anti-intellectualism', or, in Hetherington's own terminology, 'practicalism' (see Hetherington 2011). Intellectualism, succinctly put, holds that knowledge-how is a kind of knowledge-that, while Hetherington's radical anti-intellectualism argues quite the opposite: knowledge-that is a kind of knowledge-how. For example, in his 2001 book, Hetherington writes that 'Knowing-how can be better or worse; that is clear. Knowing-that is a kind of knowing-how; that is also true, even if less obviously so' (2001: 12). The two versions of Hetherington's epistemic gradualism will be further analysed in Chapter 3. At this stage, we can summarise some traits of his gradualist position:

- (i) Hetherington is a qualitative epistemic gradualist rather than a(n) (explicit) quantitative epistemic gradualist;
- (ii) Hetherington is a beyond-threshold epistemic gradualist, rather than a threshold epistemic gradualist;
- (iii) Hetherington used to be an anti-justificationist, while he seemed to have weakened his radical anti-justificationist view later;
- (iv) Hetherington's gradualism is closely connected to his theory of knowledge-how/how-knowledge.

Correspondingly, the gradualism that I will propose in this thesis is alike to that of Hetherington in regard to (i). That is to say, I will be defending the qualitative epistemic gradualism rather than the quantitative one. However, my gradualism is *different* from Hetherington's in terms of points as follow:

Firstly, I will defend both beyond-threshold and threshold gradualism (in contrast to (ii));

Secondly, I do not accept Hetherington's anti-justificationism and would *not* develop my account of gradualism on the premise that justification is unnecessary for knowledge (in contrast to (iii));

Thirdly, in contrast to (iv), I do not accept Hetherington's radical anti-intellectualism and would develop my gradualism without appealing to any specific theory of knowledge-how or how-knowledge (in particular, I will neither conclude that knowledge-how can be exhaustively reduced to knowledge-that, nor that knowledge-that can be exhaustively reduced to knowledge-how).

Now we have roughly gone through voices that support absolutism and those supporting gradualism in the current literature respectively. These can provide us with a brief sketch of the extant picture of debate around the gradability of propositional knowledge. Nevertheless, do people's attitudes towards the current debate actually stand in such a sharp contrast?

2. Minimal Gradualism and the Equivocal Intuition

As stated before, those voices advocating gradualism are just the minority. Most epistemologists have never explicitly or implicitly argued for gradualism. Perhaps it is this impression that motivates Hetherington to describe absolutism as a commonly-granted 'dogma' and conclude that 'the standard acceptance of knowledge-absolutism remains, for most epistemologists, an unshakeable assumption' (2005:130). However, one phenomenon that he might overlook is: despite the rife intuition for absolutism, most epistemologists *should* at least readily be able to be in sympathy with (if their

sympathy has not been explicitly uttered²³) ‘gradualism’ in some peculiar senses, e.g., an item of knowledge can be better justified than another one. For example, it was pointed out that:

‘It is therefore rather uncontroversial that some items of knowledge can be better justified than others. And if one wants to say that knowledge that involves better justification is better knowledge, then that strikes me as rather benign. I do not see what exactly denying this would imply.’ (Feldman 2002: n.s.)

This remark is echoed by Romy Jaster, who shares the impression that knowledge *can* be graded in a *loose* sense (e.g., in terms of the justificatory strength), but strictly speaking, ‘knows’ is not a gradable predicate:

‘We can say that the property of being tall enough is satisfied to a higher degree the taller the subject is. And we can say that the property of being sufficiently flat is satisfied to a higher degree the flatter the subject is. In the same *derived* sense we can grade knowledge relations. We can say that they are realized to a higher degree the stronger the required epistemic position is. But we should keep in mind that this is just a very loose formulation of what is really claimed.’ (Jaster 2013: 320-321)

Hetherington notes that most epistemologists take justification to be necessary for knowledge and thus reject the minimalism that he advocates. This is true. He also holds that opponents of minimalism deny that better justification yields better knowledge, and thus they would be unable to address the phenomenon of epistemic improvement. This point is called into question by Brower (2004):

‘In any case, the core problem at the intuitive level is that epistemologists who reject minimalism may grant that knowledge can be better or worse from an epistemic perspective.’ (Brower 2004: 107)

²³ This is understandable that few epistemologists have articulated their support for *internal gradualism*, as issues regarding the gradability of knowledge are seldom debated in the literature. Nevertheless, as we will see later, there are still some passages sympathetic to gradualism in the ‘internal’ sense.

In presenting internal gradualism as unorthodox, Hetherington implies that absolutists are committed to an implausible view that the quality of one's knowledge is fixed—it cannot be improved by better justification once one's epistemic position starts counting as 'knows'. In other words, one's knowledge that p can only have one possible epistemic quality. However, it is unclear why absolutists have to share this commitment. In replying to Hetherington's criticism, Adam Leite delivers a similar point:

'Hetherington is wrong to charge me with sharing the skeptic's commitment to "qualitative absolutism", the view that "all possible instances of knowledge that p share only one possible epistemic quality ... as cases of knowledge that p " (p. [308]). For my purposes here (and in my original paper) I don't have to share this commitment'. (Leite 2006: 317)

In fact, I think absolutists have reasonable grounds to reject this commitment. They could accept that the quality of knowledge can be improved by better justification just like the quality of a 'home run' can be improved by, say, a larger distance the ball travels. Nonetheless, 'hitting a home run' is still an absolute yes-or-no affair, and so is 'knowing'. Feldman, Jaster, Brower, and Leite are all sceptical about the idea that knowledge is ultimately a gradable concept. However, while sticking to their absolutist position, they are glad to admit that knowledge can be better or worse in what Hetherington calls the 'internal' sense. Their somewhat hybrid and equivocal attitude displays a natural option for absolutists to swiftly evade Hetherington's charge (*i.e.*, that absolutism fails to accommodate the idea that each part of one's justification contributes to one's relevant knowledge, as introduced in section 1.1). The idea is, knowledge can be improved in the peculiar sense that better justification yields better knowledge, but knowledge itself is *ultimately* ungradable. Just as Feldman notes, it is unclear why absolutists have to discard this option and thereby succumb to Hetherington's charge.

Admittedly, not many philosophers have reported their attitudes towards the gradability of knowledge, especially towards the question of whether a piece of better-justified knowledge counts as a better piece of knowledge. So it might be improper to assert that epistemologists *already* accept internal gradualism. However, *pace* Hetherington, I think it will be reasonable to predict that Feldman's (as well as Jaster's, Brower's, and Leite's) attitude is representative. And thus internal gradualism, at least in a loose sense, might not be as unorthodox as Hetherington takes it. It would not be surprising if most people's attitudes towards absolutism and gradualism turn out to be ambiguous, equivocal and, to some extent, inconsistent. That is because, as far as I can tell, most epistemologists more or less share intuitions as follows:

(I1) One cannot properly say that X knows that *p* better than Y does.

(I2) There is an absolute threshold for knowledge, especially, the condition of truth is absolute.

(I3) One cannot know that *p* more or less.

(I4) The strength of one's epistemic support (evidence, justification, reliability, fallibility, *etc.*) for knowledge-that-*p* can be gradable.

Those intuitions could lead to different verdicts about whether knowledge is gradable. Accordingly, people might be committed to these claims:

(I1') In light of (I1), propositional knowledge is ungradable.

(I2') In light of (I2), propositional knowledge is ungradable.

(I3') In light of (I3), propositional knowledge is ungradable.

(I4') In light of (I4), X's knowledge-that-*p* can be better than Y's (in that sense).

Being committed to (I4') means accepting a form of *minimal gradualism* in the weakest sense. However, accepting (I1'), (I2') and (I3') means being committed to some certain sorts of absolutism as well. This invites a seemingly self-conflicting attitude about the gradability of propositional knowledge. Analysing this phenomenon more intensively, we can find that people's attitude can be further subdivided:

[Attitude 1]

Some people hold that the absolutism that they accept is incompatible with the minimal gradualism, so they do not even accept the minimal gradualism and turn to reject (I4') or even (I4). I guess this group of epistemologists (if any) should be the minority.

[Attitude 2]

Alternatively, people may think (I4') *per se* is *not* incompatible with absolutism. But they cannot tolerate the equivocation to the effect that 'knowledge is gradable in this sense but ungradable in that sense'—they are keen on a 'final judgment' of the gradability of knowledge. Therefore, even though they could accept the minimal gradualism, they would eventually deny gradualism and embrace absolutism on balance. For them, the minimal gradualism cannot suffice to derive a persuasive acceptance of epistemic gradualism

simpliciter. Especially, it cannot lead to the acceptance of threshold epistemic gradualism. People who hold this attitude might believe that linguistic evidence can reflect truths about the nature of epistemological concepts—this linguistic methodology is broadly adopted in the discussion around contextualism and pragmatic encroachment. People who value the linguistic methodology much more than the significance and the force of persuasion of (I4') may accept this attitude. In addition, as we will see in what follows, (I1), (I2), and (I3) each represents a mainstream motivation for embracing absolutism. People who value these three intuitions more than (I4) might also be committed to an overall absolutist stance.

[Attitude 3]

The third kind of attitude does not treat (I4') as incompatible with absolutism, neither. People who hold this attitude do not even think it necessary to have a 'final judgment' of whether knowledge—that is gradable or not. But they just accept the minimal sense of gradualism in virtue of (I4')—beyond this, they are not open to any other form of gradualism. They are committed to (I1'), (I2'), and (I3') as well. Hence they are actually arguing for an equivocal hybrid stance: they are absolutists—in most cases; but they are gradualists too—just in the minimal sense.

Here is *my* stance: in this thesis, I will argue against all the three attitudes introduced above. Precisely put, firstly, I will certainly defend (I4') and thereby the minimal gradualism. Moreover, I will argue that (I4) is not incompatible with (I1), (I2), or (I3). So I am against [Attitude 1].

Secondly, I am in opposition to [Attitude 2] as I will not only defend the minimal gradualism. Instead, I will defend a stronger gradualism by rejecting (I1'), (I2'), and (I3'). I will reject (I1') by proving that (I1) only holds in some languages, while in some other languages, it doesn't. In addition, I will argue that the linguistic methodology is problematic so it cannot suffice to entail an overall-absolutism (see ch.2). I will reject (I2') by directly rejecting (I2) by arguing that knowledge should be seen as a *spectrum concept* that does not have a threshold at all (see ch.4). Moreover, (I3') will be denied by making the distinction between quantitative absolutism and qualitative absolutism. I will block the entailment from (I3) to (I3') by cutting off the connection between quantitative and qualitative absolutism (see ch.2).

Thirdly, as opposed to [Attitude 3], I will debunk the equivocal stance on the gradability of propositional knowledge. [Attitude 3] accepts merely the minimal gradualism, but I will propose a stronger form of gradualism. [Attitude 3] grants (I1') (I2') and (I3'), while I reject all of the three claims. Accordingly, I will demonstrate that the equivocal stance is ill-grounded.

Now it should be clear in what sense I am going to argue for gradualism and against absolutism. As we have noted before, absolutism has been seen as the orthodoxy for a long time. So why is absolutism so appealing? What are motivations for accepting absolutism? The following section will provide three representative arguments for epistemic absolutism.

3. Arguments for Epistemic Absolutism

The idea that knowledge-that is ungradable is deeply entrenched in standard epistemological thinking. Nevertheless, few philosophers have seriously developed systematic arguments for epistemic absolutism, which causes Hetherington to call absolutism a 'dogmatism'. Although there are few systematic arguments for absolutism in the history of epistemology, which

makes absolutism dogmatic to some extent, some motivations for absolutism can still be summarised or conceived. Here, three typical ways to argue for epistemic absolutism will be given.

3.1: Argument from Linguistic Evidence

The most straightforward motivation for people to reject gradualism might be various linguistic evidence in our ordinary language revealing the oddness of using 'know' as a gradable verb, just like what is shown at the beginning of this chapter. Those linguistic data displayed before have to some extent revealed that epistemic gradualism goes against our linguistic intuition (at least in English). Hetherington (2001) starts his defence of gradualism by listing several expressions involving *gradable* uses of 'knows that' that he takes to be consistent. For example:

'I know very well that I have hands; I know less well that I have a brain. (A neurosurgeon could know much better than I do that I have a brain. I do not want such good knowledge of my having a brain!)

'I know that I feel pain. I know that you do, too. But I know better that I do than that you do—whereas you know better that you do than that I do!

'I know fairly well that the earth is round, although many other people know this much better than I do.' (2001: 1-2)

Hetherington argues that these gradable uses of 'know that' can serve as counterevidence to the linguistic support for absolutism. To summarise, these expressions involve two indicators of gradability: degree modifiers such as 'very well' or 'fairly well', and comparative constructions such as 'know better that' or 'know less well that'. Although Hetherington's linguistic objection to absolutism barely received *direct* challenges, we can still find competing arguments in the literature which could be seen as *indirect* but unneglectable replies to him on behalf of absolutists. In what follows, I will introduce two linguistic arguments for absolutism from Stanley (2005) and Dutant (2007)

respectively, and illustrate how these arguments can threaten Hetherington's alleged linguistic counterexamples.

Jason Stanley

In the second chapter of his 2005 book, Stanley famously argued that knowledge ascription is not gradable. By comparing the uses of 'knows' in our daily language that those of other paradigmatic gradable adjectives or verbs (e.g., 'tall' 'flat' and 'like' 'regret'), Stanley concluded that 'knows' cannot pass the two linguistic tests for being a gradable predicate: 1) allowing for degree modifiers like 'very' and 'really'; 2) being conceptually related to felicitous comparative constructions. Although the opponents that Stanley aimed to defeat were contextualists such as Cohen and DeRose, it is noteworthy that the way Stanley argued for absolutism was exactly diametrically opposed to Hetherington's defence of gradualism.

Firstly, the degree modifiers test. As Hetherington (2001) points out, there are *prima facie* consistent statements like 'John knows *very well* that Trump is the president', which apparently indicates that 'knows' allow for modifiers. Stanley denies that the phrase 'very well' here is modifying the knowledge relation. The idea is, if it is 'knows' that is felicitously modified by 'very well', then the statement's negation should also be felicitous. However, by comparing expressions like 'John does not *know* very well that Trump is president' with expressions like 'John does not *like* Trump very much', Stanley argues that while the latter sounds natural, the former turns out to be untoward. This indicates that 'knows' significantly differs from typical gradable predicates such as 'likes'. Stanley thus concludes that the phrase 'very much' is not a way to modify 'knows that', but rather a pragmatic indicator similar to its occurrence in sentences like '2 is very much an even number' (2005: 38). In addition to 'very well', modifiers such as 'really' are also examined by Stanley. For contextualists, it seems to be natural to utter claims like 'I guess I don't really know that the bank is open', where 'know' naturally occurs with 'really'. Stanley rejects these *prima facie* cases for

gradualism by arguing that the negations of modified knowledge ascriptions cannot consistently be conjoined with assertions of the unmodified forms. For example, it seems self-conflicting to say that 'S knows that p , but does not really know that p '. On the contrary, it is fairly consistent to say 'S is tall, but not really tall' (see 2005: 37). Stanley's diagnosis of this asymmetry is straightforward: 'tall' is gradable, while 'knows' is not. The use of 'really' in knowledge ascriptions, for Stanley, is just a *hedge*, of which the linguistic function 'is to comment on the appropriateness of asserting the embedded sentence ... in so using "really", one concedes the infelicity of asserting that one knows the proposition in question.' (Stanley 2005: 46)

When it comes to the test of comparative constructions, Stanley also denies that 'knows' can naturally occur with comparative constructions. Stanley does not agree that idiomatic constructions such as 'S knows that p better than anyone' would suffice to serve as proof that knowledge ascriptions are *genuinely* gradable (see Stanley 2005: 40). Rather, he argues that sentences like 'That broccoli is low fat is better known than that broccoli prevents cancer' (ibid: 41) does not genuinely show that 'knows' can be conceptually related to a qualitative comparative construction, because the use of 'better known' in this sentence does not imply that the fact that broccoli is known in a better way or with better justification, but just that the fact is known by *more people*. This diagnosis can be supported by the oddness of the following statement: 'It is well known that p , and less well known that q , but more people know that q than know that p .' (2005: 42) If it is only the amount of 'knowers' that is put into comparison, then those comparative constructions involving 'knows that' neither show that knowledge-that is qualitatively gradable, nor prove it to be quantitatively gradable.

Julien Dutant

Another representative argument for absolutism from linguistic evidence is given by Julien Dutant, who starts his argument with some objections to Stanley. Following Stanley, Dutant (2007) firstly examines the gradability of

'knows' in terms of the two linguistic criteria for being gradable predicates, namely, 'degree modifiers and comparative constructions'. Stanley concludes that 'knows' cannot pass the test of degree modifiers, while Dutant objects that there are natural statements where 'knows' is genuinely modified by degree modifiers²⁴. For example:

- 'a. Tim knows well that we must have no illusions that somehow the danger has passed. (Dick Cheney)
- b. But Chirac knows well that, while times change, an unchanging rule of politics is this: you say whatever you have to say to get elected. (International Herald Tribune)
- c. And he knows very well that it requires work to get past old grievances. (GW Bush)' (2007: 5)

Dutant argues that the uses of 'well/very well' in these statements should not be understood as pragmatic indicators that confirm that the subject knows. Instead, they do imply that the subject's knowledge enjoys higher quality. In response to Stanley's criticism that those *prima facie* felicitous modified knowledge ascriptions would turn out to be awkward once they are negated, Dutant points out that the same alleged problem also applies to paradigmatic gradable predicates, such as 'shows' and 'sees'²⁵. For instance:

- 'a. You see very well that he is lying.
- b. ?You do not see very well that he is lying.
- a. The results show very well that the company fares better.

²⁴ Hetherington (2011 §5.13) provides a different response to Stanley's objection on the basis of his *practicalism, i.e.*, an account of knowledge that knowledge is essentially a kind of ability rather than a kind of belief. In response to Stanley's test of degree modifier, Hetherington concedes that 'very much' is not a way to modify the knowledge relation, but he insists that modifiers like 'very well' can understandably and genuinely modify 'knows' as long as we understand knowledge as a kind of ability.

²⁵ Hetherington (2011 §5.13) advocates a more steadfast reply regarding this problem of negation. He argues that we do understand utterances such as 'He knows, but not as well as he should, that he is now a man.' That utterance can make sense as it is conceivable that 'Maybe he has enough evidence and an unconfident belief for the knowledge; yet he could easily permit his lack of confidence to unseat his belief if even slight counter-evidence appears' (2011: 210). Even though I also appreciate Hetherington's argument, I think Dutant's reply might be more undercutting than Hetherington's, as it relies less on the readers' linguistic intuition.

b. 'The results show very well that the company fares better.' (2007: 9-10)

Therefore, the oddness of negations of 'knows well that' does not suffice to disprove the gradability of 'knows'. In summary, it is unfair to claim that 'knows' does not allow for degree modifiers because of linguistic data.

As for the test of comparative constructions, Dutant provides, *contra* Stanley, many cases (see Dutant 2007: 7 & 25) where 'knows-that' constructions can be *genuinely* and felicitously built with qualitative comparatives. For example:

(17) Myers **knows better than** Mr. Kennedy that our military will have 'the will to win.' (Letter to the Washington Times)

(18) It works, but Julie **knows better than** Sarah that she cannot let Frank off the hook. (Alan C. Shaw, MIT Prof.)

(19) Mike Matthews **knows even better than** Norton that durability is the name of the game for the lefty reliever. (Cincinnati Post) (2007: 7)²⁶

For Dutant, it seems that the two linguistic tests cannot constitute reasons for denying the gradability of knows, and thus he carves out a new way of defending absolutism. Drawing on Kennedy and McNally (2005), Dutant employs the so-called 'two-scales account' of quantitative/qualitative modifiers to lend support to absolutism. According to the two-scales account, the scale over which 'know well/better' (qualitative degree modifications of 'know') constructions are evaluated is orthogonal to the scale over which 'know much/a lot' (quantitative degree modifications of 'know') constructions

²⁶ Hetherington (2011 §5.13) adopts a similar strategy to address Stanley's test of comparative constructions. That is, Hetherington argues that there are linguistically acceptable utterances involving 'knows better than' constructions as long as there is *evidential room* for one to know better or worse, for example, 'John knows better that he is in pain than Mary does' (2011: 212) and 'That broccoli is a vegetable is better known (by all who do know it) than is its being low in fat (by all who know this)' (2011: 213). Both Hetherington's reply and Dutant's reply have their own merits. Dutant's objection is appreciable as it quotes real English expressions from authoritative sources. Hetherington's argument is helpful in the sense that it explains more clearly *why* some knows-relevant comparative constructions are linguistically acceptable while some are seemingly less felicitous—the difference lies in whether the utterance itself involve sufficient conceivable *evidential room* for one to know the proposition in question better or worse.

are evaluated. Although Dutant discovers that 'knows' can be modified by genuine qualitative degree modifiers like 'well', it was argued that qualitative modifications of 'know that' does not actually denote degrees of knowledge. His argument goes as follows:

'Know well that' constructions *do* have degrees, but the problem is whether this output of the well-modification, *qua* a new gradable predicate, operates on the same scale as the original unmodified 'know that' construction. For Dutant, the way in which 'know' is modified by 'well' is analogous to the way in which 'written' is modified by 'well'. Dutant argues that 'written' itself is associated with the *quantitative* scale and has denoted a maximal point of that scale. The idea is, when one says that 'the article is written', one means that the article is *completely* written—if the article is just half-written, it is not genuinely written. Because the unmodified 'written' itself has already denoted the maximal endpoint of the quantitative scale, the 'well written' modification cannot be understood as denoting a higher degree along the same scale. That is to say, the well-modification can only operate on the *qualitative* scale, which is orthogonal to the scale which the unmodified 'written' is associated with. Call this *the two-scales argument*. Dutant then applies the same argument to 'know'. Accordingly, it is concluded that the scale over which 'know well that' is evaluated is orthogonal to (and thus different from) the scale over which the unmodified 'know that' is evaluated. Now that when we claim that 'S knows well that p', we mean that S's knowledge is of high quality, it seems that the well-modification of 'know' is corresponding to the qualitative scale. Given the two-scales account, the relevant scale over which the *unmodified* 'know that' construction is evaluated must be orthogonal to that of its well-modification. Therefore, if 'know that' can be graded, then it can only be graded along the quantitative scale. However, because 'know that' cannot felicitously occur with quantitative degree modifications—for example, it is awkward to say 'Bob knows a lot/much that the city is unsafe' (Dutant 2007:12)—'know that' cannot be properly evaluated along the quantitative scale, either.

In summary, according to Dutant, ‘know that’ is *not* evaluated along the qualitative scale, and *cannot* be felicitously evaluated along the quantitative scale, therefore, ‘know that’ cannot be graded along any scale. So, if knowledge-that does not admit of degrees, then why are some ‘know well that’ modifications felicitous? Dutant’s explanation is, their apparent felicitousness comes from the widely-accepted gradability of justification—the ‘know well that’ modification only denotes degrees of justification rather than degrees of knowledge.

3.2: Argument from Object

The second form of argument for absolutism might be motivated by the well-known T-schema, *i.e.*, ‘(for all p) $\langle p \rangle$ is true iff p ’ (see Båve 2013), that is, a proposition p is true if and only if p . Therefore, asserting a proposition p is equivalent to asserting that p is true. For example, to say that ‘snow is white’ equals to say that “‘snow is white’ is true” (see Ramsey 1927; Ayer 1935; Quine 1970; *etc.*). Trading on this intuitively appealing view, it seems natural to accordingly draw a conclusion as follows:

[Equivalence Thesis of Knowledge]

Because the proposition that p is equivalent to ‘ p is true’, knowing that p is equivalent to knowing that p is true.

The equivalence thesis of knowledge (hereafter, ETK) seems to make *prima facie* sense, if the T-schema is acceptable—the thesis just involves a simple substitution given that p and ‘ p is true’ are mutual-substitutable. According to ETK, to know that snow is white is to know that ‘snow is white’ is true; to know that Trump is the president equals to know that ‘Trump is the president’ is true. With this equivalence relation in play, it is seemingly smooth to derive

that to know a proposition is to know the truth value of that proposition. That is to say:

[Object of Knowledge Thesis²⁷]

The *object* of a piece of proposition knowledge is the truth of that proposition.

Taking one step further, one can argue that to have a better knowledge of a proposition is to have better knowledge of the truth value of that proposition on the basis of OKT. Absolutists can then argue that, however, the fact that one proposition is true is absolute. If a proposition is true, it is completely and perfectly true—the cut-off point between a true proposition and a false proposition is absolute and clear, no true proposition is ‘truer’ than another proposition. Therefore, knowing that Trump is the president requires nothing more than to know the truth of the fact that Trump is the president. Once this condition (knowing that ‘Trump is the president’ is true) is met, any further justification or any more evidence cannot or any more solid belief toward that proposition cannot contribute more to the truth of the fact that Trump is the president. Even working in the White House cannot improve your knowledge that Trump is the president is true, if you have already known that Trump won the US selection by reading newspapers before you enter the White House and work with President Trump. The truth of ‘Trump is the president’ has already be known by you, the threshold has already been crossed. The fact that ‘Trump is the president’ is perfect as a fact—if it is true, then it is perfectly true. Therefore the knowledge of that fact, once achieved, is perfect as a piece of knowledge of that fact—no room for improvement, just absolute and perfect knowledge. In light of this, propositional knowledge is absolute.

To clarify the schema of this form of argument for absolutism, the argument can be summarised as follows:

²⁷ Abbr., OKT.

[The Argument from Object]

(AO1) Asserting a proposition p is equivalent to asserting that ' p is true', therefore knowing that p is equivalent to knowing that p is true. (ETK)

(AO2) To know a proposition is to know the truth of that proposition. (OKT)

(AO3) The truth of a proposition is absolute.

(AOC) Knowledge of a proposition is absolute.

This Argument from Object captures the sense of what Hetherington (2005) calls as 'the strongest form of defence' for epistemic absolutism:

'To know that p is to know the truth of the proposition p ; and you cannot alter, even slightly, the object of that knowledge without wholly deleting the knowledge that p . (Either you lose knowledge altogether, or you replace knowledge that p with knowledge that q , where $q \neq p$.) Knowing that p is an all-or-nothing state of affairs. If any different component of the knowledge that p (such as the justification) is altered, even by being improved, again that particular piece of knowledge that p is no more. At best, you gain a new instance of knowledge that p ²⁸... Contrast knowing that p with knowing a person, say. There can be better or worse knowledge of a person—because that kind of object of knowledge has more or less of itself being known at a given time.'
(Hetherington 2005: 148)

²⁸ Surely, there is a trivial understanding of Hetherington's 'If any different component of the knowledge that p (such as the justification) is altered, even by being improved, again that particular piece of knowledge that p is no more. At best, you gain a new instance of knowledge that p ', which interprets this quotation as saying that once an instance of knowledge-that- p is improved, it is no longer the same original instance of knowledge-that- p -- simply because it becomes a better instance of knowledge-that- p ! This interpretation is too trivial to support absolutism, and seems endorse that there can be 'better knowledge that p ' (which is just what gradualism aims to urge), but just argues that it is not the same instance of knowledge anymore (certainly, every improved item is not exactly the same item in this sense). In addition, this reasoning seems to apply to knowledge of person as well. It makes perfect sense to argue that Mary's improved knowledge of John is not her original knowledge of John just because it is improved. If that is the case, this would fail to retain the disanalogy between knowledge of proposition and knowledge of person that Hetherington presented in that cited paragraph. Therefore, I would not discuss this trivial interpretation here.

While conceiving such a defence, Hetherington emphasises that it might be argued that the object of knowledge-that- p is the truth of the given proposition p . This object cannot be more or less known or known in part, which is different from the object of knowledge of a person (*i.e.*, that person *per se*), because a person can be known more or less at a given time. Likewise, this object cannot be altered or improved without changing the original instance of knowledge-that- p into an instance of knowledge of another proposition q , because if more evidence of p is gained, then the subject just gains a new piece of knowledge of the new evidence. For example, it might be suggested that one's knowledge that 'Trump is the president' can be improved by gaining more evidence proving that fact, *e.g.*, by seeing Trump giving a presidential speech in the television. However, according to the Argument from Object, seeing Trump giving a presidential speech in the television does not improve your knowledge of the truth of 'Trump is the president', instead, it at best just provides with a new piece of knowledge of a new proposition that 'Trump is giving a presidential speech in the television'. The truth of 'Trump is the president' cannot be improved by the fact that 'Trump is giving a presidential speech in the television'.

This form of argument for absolutism can be utilised to defend both quantitative epistemic absolutism and qualitative epistemic absolutism. On one hand, it can be argued that this argument capture perfectly the sense, as what Ryle's remark conveys, that a subject 'could not be said to have incomplete knowledge of Sussex being an English county. Either he knows this fact or he does not know it.' (Ryle 1949: 59) Because the truth of the fact that 'Sussex is an English county' cannot be known more or less or in part—if it is known, it is known completely and perfectly—then according to the argument from truth value, the knowledge of Sussex being an English county cannot be quantitatively gradable. On the other hand, it can be interpreted that the truth of a proposition is absolute means that a truth cannot be qualitatively better than another truth in terms of qualifying as a truth. This can explain what Dretske said that 'factual knowledge is *absolute*. It is like

being pregnant: an all or nothing affair' (1981: 363). The truth of a proposition is an all or nothing affair—the threshold is clear-cut and absolute. Once the threshold is crossed, every truth is equal *qua* a truth. Neither more evidence nor better justification can improve the truth and thereby the corresponding factual knowledge. In that propositional knowledge is not qualitatively gradable.

In addition, it seems that absolutists can trade on this argument to explain the linguistic evidence supporting absolutism. It can be argued that when asserting that someone knows that *p*, we are in fact asserting that he or she knows that *p* is true. Accordingly, when evaluating one's knowledge of a certain proposition, we are actually evaluating her knowledge of the object of that instance of knowledge-that-*p*. The oddness of asserting 'S knows that *p* better than another person' can be explained by the oddness of asserting 'the truth of *p* is better than the truth of *q*' as the truth of *p* is the object of knowledge-that-*p*. Those *prima facie* linguistic cases for epistemic gradualism can then be explained as just cases where the object of the original instance of knowledge is alerted. For example, 'Myers knows better than Mr. Kennedy that our military will have "the will to win"' can be explained as just implies that Myers knows more *relevant facts* than Mr. Kennedy about the fact that the military has the will to win, but these relevant facts just entitle Myers of more knowledge of relevant evidence, rather than genuinely make Myers' knowledge of the fact 'our military will have "the will to win"' a better instance of knowledge than Mr. Kennedy's. The linguistic explanation derived from the Argument of Object is different from and even incompatible with Dutant's. Dutant holds that 'S knows well / better that *p* if and only if one knows that *p* on the basis of a good number of / more propositions' (Dutant 2007: 16), while the Argument from Object will argue that by knowing a good number of / more propositions, one just alerts the object her original knowledge-that-*p*, rather than knows better that *p*. In addition, Dutant wrote that 'Working in the White House does not improve one's knowledge that G.W. Bush is president. By contrast, having been a women basketball coach

for twenty years improves one's knowledge of the fact that women basketball has come a long way' (2007: 8). That is to say, Dutant endorses that sometimes a certain instance of knowledge-that-p can be improved by gaining more supporting fact.

3.3: Argument from Contextualism

The last motivation for accepting epistemic absolutism that I want to introduce is more like an impression or an induction rather than a rigid argument. This impression arises from the phenomenon that debates around epistemic gradualism and absolutism are always relevant to (perhaps originate from) the debate around contextualism and anti-contextualism. Especially in virtue of the wide influence of Stanley's discussion about gradualism, people may be inclined to combine gradualism with contextualism, and believe that by rejecting gradualism, contextualism can thereby be undermined. It is thus natural for opponents of contextualism to be averse to gradualism. One may worry that if gradualism cannot be independent of contextualism, especially if gradualism turns out to be rooted in contextualism, then because contextualism is implausible, gradualism should thus be rejected. This worry is reasonable, especially when taking the (T1) thesis of threshold epistemic gradualism into consideration—denying that there is a fixed ungradable standard of knowledge sounds extremely alike to a contextualist claim.

The argument can be formulated in a form like this:

[Argument from Contextualism]

(AC1) Epistemic gradualism is based on epistemic contextualism.

(AC2) Epistemic contextualism is wrong.

(ACC) Epistemic gradualism is wrong.

Advocates of contextualism would not trade on this argument to reject gradualism, because they would deny (AC2). However, for those people who dislikes contextualism and finds that in previous discussion around gradualism, 'knows' is always taken to be analogous to context-sensitive adjectives like 'talk', 'flat' and 'large' (in order to support contextualism), it is natural for them to inductively conclude that (AC1), and thus rejects gradualism. So gradualists are required either to defend contextualism (to disprove (AC2)), or to prove that epistemic gradualism can be motivated and developed independently of contextualism (to falsify (AC1)). In this thesis, I will take the latter approach, i.e., to reject (AC1), by displaying how gradualism can be constructed without thereby being committed to contextualism. It will also be argued that the debate about gradualism/absolutism should not be a sub-battlefield of the debate around contextualism—the debate has its *sui generis* theoretical value.

4. Concluding Remarks

This chapter has introduced two conflicting views on the gradability of propositional knowledge, *viz*, epistemic absolutism and epistemic gradualism. Different forms of epistemic gradualism have been characterised, and I have clarified the form of gradualism that this thesis aims to defend. Three main arguments for epistemic absolutism have been introduced. In the next chapter, I will argue against all the three forms of argument for epistemic absolutism introduced above respectively. It will be concluded in that chapter that there is no plausible reason to accept epistemic absolutism. Now that absolutism is ill-grounded, it should do no harm to turn to reassess gradualism.

Chapter Two: Why Absolutism is Ill-Founded

Abstract: This chapter purports to demonstrate that epistemic absolutism is not as well supported by arguments as the orthodoxy has it. The first section will debunk the linguistic argument for absolutism, which appeals to the alleged fact that we usually do not use 'knows' as a gradable predicate in the ordinary language. Drawing on Hetherington's criticism of absolutism, I will develop a more thorough objection to the linguistic argument for absolutism. My objection consists of two parts. As Stanley and Dutant, on absolutists' behalf, have put forward two rejoinders to Hetherington's original objection, the first part of my objection will focus on refuting the two rejoinders respectively. After that, the second part of my objection will propose three general strategies that gradualists can take to reject all forms of linguistic argument for absolutism. I will demonstrate that: 1) gradualism can also provide counterevidence supporting the gradability of knowledge-that in the English language; 2) non-English linguistic data can also lend support to gradualism; 3) the linguistic argument for absolutism is premised on a linguistic methodology that is untenable because of the gap between an epistemic concept's linguistic uses and its conceptual nature. The second section puts weight on the argument from object. I will reject one essential premise of this argument, to wit, 'to know that p is to know that p is true'. Besides, it will be argued that the absoluteness of knowledge's object cannot entail the absoluteness of knowledge per se. Finally, the third section will briefly show how gradualism can be constructed without appealing to contextualism. The conclusion of this chapter is: epistemic absolutism is ill-founded.

1. Objections to the Argument from Linguistic Evidence

1.1: Objections to Dutant

The most influential and straightforward reason to reject gradualism might be the linguistic evidence that 'knows' seems not to be a gradable verb in English. In the first chapter, I have outlined the two most influential and systematic linguistic arguments for absolutism. To reject absolutism, gradualists have to refute the two arguments put forward by Stanley and Dutant. A (somewhat) good news for gradualists is that Dutant has provided some cogent criticisms to Stanley which I will adopt at this stage. As we have noted in the last chapter, to sum up, Dutant's objections consist of two parts. First, Stanley neglects that there are also cases where 'know that' can be felicitously modified by degree modifiers and be built with comparative constructions. Second, Stanley argues that, unlike paradigmatic gradable terms such as 'regrets', 'know well that' can hardly be infelicitously negated. Notwithstanding, he ignores that there are also paradigmatic gradable verbs like 'shows' and 'sees' whose negation or question constructions are infelicitous, hence the oddness of negation and question construction of 'know well that' fails to prove its ungradability.

So, now there remains Dutant's own argument. Dutant's argument for absolutism essentially hinges on a quickly drawn analogy between 'know that' and 'written'. It is presupposed that the two-scales account, which is alleged to be applicable to 'written', also holds for 'know that'. However, it is not explained why 'know that' is analogous to 'written' rather than some gradable verbs such as 'like' or 'enjoy'. This makes Dutant's argument suspicious as it seems that he implicitly presumes, rather than proves, that 'knows' is ungradable. Apart from this, the analogy between 'know that' and 'written' is improper as there is a significant asymmetry between the two phrases such that *the two-scales argument*, which is applicable to the latter, is inapplicable to the former. Dutant's two-scales argument, which argues that the scales along which 'well written' and 'written' are measured have to be orthogonal, turns essentially on two premises:

(P1) 'Written' has denoted the maximal point of the quantitative scale.

(P2) 'Written' *can* be evaluated among the quantitative scale.

Note that, (P1) is premised on (P2)—after all, if a term cannot be evaluated along the quantitative scale, how can it denote the maximal point of the quantitative scale? 'Written' can be evaluated along the quantitative scale because it is natural to say 'Bob has written a lot of papers' (although for every single paper that is written, it is *completely* written). Therefore, if 'know that' is sufficiently analogous to 'written' such that the same two-scales argument also applies to 'know that', then 'know that' should also meet two conditions:

(P1*) 'Know that' has denoted the maximal point of the *quantitative* scale.

(P2*) 'Know that' can be evaluated among the quantitative scale.

However, this is not true. Just as Dutant admits, 'know that' does not admit of quantitative degrees, so it is also inconsistent to claim that 'know that' construction denotes the maximal point of the quantitative scale. Therefore, the analogy between 'know that' and 'written' is improperly asymmetric, and it is unclear why the scale along which 'know well that' is evaluated has to be orthogonal to the scale along which 'know that' is evaluated. As a matter of fact, Dutant fails to rule out the possibility that 'know that', as well as 'know well that', can be evaluated among the qualitative scale. The dilemma is thus disarmed.

So why does Dutant think that the analogy between 'written' and 'know that' holds water? That seems to be because he presumes (without demonstration) that the two-scales account put forward by Kennedy and

McNally, which focuses on adjectives, also applies to (all) verbs. When analysing the two-scales account, Dutant notes that:

'The account has two crucial features. First, the two scales are in principle orthogonal. The scale over which one measures *how much* of the thesis is written is orthogonal to the scale over which one measures *how well* the thesis is written. Second, the relevant scale for applying the basic predicate is the quantitative scale (*very / much*), not the qualitative one.

The same points seem to hold for verbs. The scale over which we measure whether John works a lot is orthogonal to the scale over which we measure whether he works well. And the relevant scale for saying whether he worked or not is the quantitative one.' (2007: 12)

But this presumption is made too quickly, as the two-scales account does not always hold for verbs. Many counterexamples can be found, for instance, 'understand that'. The 'understand well that' modification is associated with the qualitative scale, e.g., 'Sam understands well that he cannot force others to love him'. But this does not imply that the relevant scale for saying whether Sam understands that *p* is the quantitative one. It is even infelicitous to claim that one understands a lot that *p*, in that 'understands' is akin to 'knows' in terms of the lack of quantitative gradability. Nonetheless, 'understand that' is still a paradigmatic qualitatively gradable term, to which the two-scales account fails to apply. Besides, the qualitative 'design-well' modification (e.g., 'Ludwig designed this building well') is evaluated among the qualitative scale, but it does not mean that the only relevant scale for measure whether Ludwig designed this building is the quantitative one. Quite the opposite, it seems to be awkward to say something like 'Ludwig designed this building a lot'.

Now we have seen why both Stanley's and Dutant's arguments fail. However, it is completely conceivable that absolutists can argue that their arguments fail to represent the best shot that absolutism can give. Also, one might find the above gradualist objections to absolutism too particular to be of broader persuasiveness. In light of this, the remainder of this section will

explore three more general strategies that gradualists can take to refute all forms of linguistic argument for absolutism.

1.2: Objection from Counterevidence

Arguments for absolutism from linguistic evidence are typically based on cases where our gradable uses of the 'know that' construction are odd. However, competing linguistic evidence supporting epistemic gradualism can also be found. As we have noted before, both Hetherington (2001) and Dutant (2007) provide relevant counterevidence indicating that 'know that' can also be felicitously used in a gradable manner. In particular, linguistic data listed by Dutant (2007) are expressions with authoritative sources that he collected from the Internet. The existence of those competing linguistic data does not imply that we *have to* use 'know that' gradably, but it does show that we *do* use it gradably sometimes.

Here, I do not plan to repeat their counterevidence or list more linguistic data supporting gradualism. I believe that counterevidences provided by Dutant and Hetherington have already been sufficient to reveal an important fact. That is, linguistic data can also be found to support epistemic gradualism in the English language. This nip-and-tuck state of play is actually in epistemic gradualism's favour, as gradualists do not need to build their view on the basis of linguistic evidence, while epistemic absolutism is primarily motivated by negative linguistic data for gradualism. Gradualism can profit from the fact that absolutists *cannot* reject gradualism on the mere basis of linguistic evidence.

Moreover, another point that is worth noting is that the awkwardness of gradable uses of 'know that' might not be able to lend as strong support to absolutism as it is expected to be. After all, for some paradigmatic gradable verbs, negative linguistic data can also be found. For example:

- (1) I very much *believe* that it will rain tomorrow.
- (2)???I believe very well that it will rain tomorrow.
- (3) I believe that it will rain tomorrow more strongly than Tom does.
- (4)???I believe better that it will rain tomorrow than Tom does.
- (5) I *remember* well that it rained last week.
- (6)???I remember very much that it rained last week.
- (7) I remember more clearly that it rained last week than Tom does.
- (8)???I remember better than Tom that it rained last week.

'Believe' and 'remember' are two typical gradable verbs. Belief and memory are also ordinarily granted as gradable in terms of the firmness, clarity, credence, *etc.* Nevertheless, there are still infelicitous gradable utterances containing these two verbs, *e.g.*, (2), (4), (6), and (8). The moral is: the existence of infelicitous gradable uses of a concept does not suffice to deny the gradability of that concept.

Nevertheless, I am far from believing that all linguistic data, no matter supporting gradualism or absolutism, are intuitive for all readers. It is understandable that linguistic intuition sometimes varies from person to person. Moreover, it can also vary from language to language. This leads us to the second anti-absolutist strategy.

1.3: Objection from Non-English Language

The extant absolutism/gradualism debate is basically limited to the English language. This should be seen as a worrisome problem, given that the relevant debate significantly hinges on the linguistic intuition, while people's

linguistic intuition can be influenced by the language that they speak. The current English-centred investigation of gradualism overlooks this influence. In some other languages, people's linguistic intuition towards gradable uses of 'know that' can be different. For example, 'A knows that p better than B' might be a bad expression in English, but there is nothing odd to say that '我比你更知道.....' in Chinese, which means 'I know that ... better than you do'. In addition, there is another synonym of 'know that' (知道) in Chinese, namely, '清楚', which means 'know clearly that', that can be used even more naturally as a gradable verb. Hence, differences in linguistic felicitousness will emerge when we compare the following utterances:

(9)???I know better that I suffered a bad toothache last night.

(10) 我比你更知道我昨晚牙很疼。

(11)???I know clearly better than you that he is a nice guy.

(12) 我比你更清楚他是个好人。

Utterances (9) and (10) deliver the same proposition in two different languages respectively - (9) is just the Chinese translation of (10). So are utterances (11) and (12). Although they share the same semantic meanings, the English expressions are infelicitous while the Chinese ones are natural. This indicates that a negative datum in one language can be a positive one in another language. One might doubt that I just reported the linguistic intuition of myself, which might not be representative enough to reflect the linguistic intuition held by a broader community of Chinese speakers. To save my argument from this suspicion, let me exhibit more linguistic data that are collected from authoritative websites or published books in the Chinese cyberspace. The original Chinese expressions and their corresponding English translations are listed below:

(13) 当我去了好学校, *更加知道*读书无用都是骗人的。

When I went to a good school, I *knew better that* it is deceiving to claim that study is useless.

(http://www.sohu.com/a/208979258_827871)

(14) 我想你比我**更知道**，你就是我的女主角。

I think that you *know better* than I do that you are my heroine²⁹.

(<https://kuaibao.qq.com/s/20180624B0VHFZ00?refer=spider>)

(15) 我只是看透了自己，知道自己所想，所欲，所惧，正视自己的阴暗，面对自己的丑陋，**更加知道**每个人都是不同的，所以可以理智的体会别人的所想，不掺杂自己的意愿。（麦琦，《贵妇深藏不露》）

I just saw through myself, knew my thought, my desire, my fear, faced up to the dark side of myself, the ugliness of myself, came to *know better that* everyone is different, and thus was able to sympathise with others more rationally without being affected by my own will.

(Maiqi, *A Undiscovered Noblewoman*; see

<https://www.juzimi.com/article/353563>)

(16) 我知道淚該停，我知道夢該醒，我**更知道**我還是不肯相信（劉志宏 & 劉思銘，《淚該停夢該醒》，歌詞）

I know that my tears should be stopped, I know that my dream should be over, but I *know better that* I am still not willing to believe this. (Zhihong Liu & Siming Liu, *The Tear Should Be Stopped, the Dream Should Be Over*, lyrics; see <https://mojim.com/twy100441x2x9.htm>)

²⁹ There is a comma in the original Chinese sentence because in the Chinese language, a comma is ordinarily inserted between the predicate and the clause in order to highlight the object clause.

The existence of these data proves that people do use ‘知道’ (*i.e.*, ‘know that’) gradably in the Chinese language. Admittedly, not all gradable expressions involving ‘know that’ constructions are felicitous in Chinese. The acceptability of each utterance following decreases progressively:

(17) 我比你更知道他是个好人 (I know better than you that he is a nice guy) 。

(18)?我比你更知道雪是白的 (I know better than you that snow is white)。

(19)??我比你更知道卡夫卡生于 1883 年 (I know better than you that Kafka was born in 1883)。

(20)???我比你更知道 $1+1=2$ (I know better than you that $1+1=2$)。

It is beyond the purpose of this chapter to provide a full explanation for this decline of acceptability. However, I believe that the oddness of a gradable use of ‘know that’ or ‘知道’ might be positively correlated with the difficulty in conceiving *how* one can know the given proposition better than another person. It is relatively easier to understand how Jim can know that Hugh is a nice person better than Tom does—probably Jim is a close friend of Hugh while Tom is not. In contrast, it is relatively hard to imagine a relevant context in which one can know that $1+1=2$ better than another person does. With these being said, it should be evident that linguistic data cannot serve as a good basis for epistemic absolutism, as the linguistic intuition behind can vary from language to language—some negative data for gradualism in English are positive in other languages.

1.4: Objection from Methodology

The two objections introduced above both focus on analysing how ‘know that’ is used in our ordinary language, while the last strategy that I plan to propose

is more undercutting. The two strategies stated before both aim to show that the linguistic data do not favour absolutism over gradualism, but they do not doubt the role that linguistic evidence can play in a debate around knowledge's conceptual nature. Hence they can be taken to be two *internal criticisms* of the linguistic argument for absolutism. In contrast, the third refutative strategy that I am about to introduce can serve as an *external criticism*. It questions the linguistic argument for absolutism by undermining its postulated methodology, *viz*, the methodology that generalises the epistemological nature of a concept from its daily uses in the ordinary language. This linguistic methodology is widely adopted in contemporary epistemology. Faith in this linguistic methodology also constitutes the primary motivation for opponents of gradualism to hold the *overall-absolutist attitude* and the *equivocal-absolutist attitude* (*cf.* Chapter One, Section 2).

Recently this methodology was challenged by Hazlett (2010, 2012) who argues that non-factive uses of 'know-that' can be found in people's ordinary linguistic practice, which contradicts the longstanding 'consensus' that propositional knowledge has to be factive. In other words, according to Hazlett's discovery, linguistic data can be at odds with the widely-accepted thesis of 'Factivity':

Factivity Certain two-place predicates, including 'knows'³⁰, 'learns', 'remembers', and 'realizes', which denote relations between persons and propositions, are factive in this sense: an utterance of 'S knows p' is true only if p, an utterance of 'S learned p' is true only if p, and so on. (Hazlett 2010: 499)

The upshot is that Factivity is falsified by our non-factive uses of 'knows'. And thus, if one chooses to stick to the linguistic methodology, then one has to give up the standard truth condition of propositional knowledge. Here are some non-factive uses of 'knows' that Hazlett provided:

³⁰ Here, I will only focus on 'knows' and set other predicates aside for the purpose of the present section.

- (1) Everyone knew that stress caused ulcers, before two Australian doctors in the early 80s proved that ulcers are actually caused by bacterial infection.
- (2) He figures anything big enough to sink the ship they're going to see in time to turn. But the ship's too big, with too small a rudder ... it can't corner worth shit. Everything he knows is wrong.
- (3) In school we learned that World War I was a war to 'make the world safe for democracy,' when it was really a war to make the world safe for the Western imperial powers.
- (4) I had trouble breathing, sharp pains in my side, several broken ribs and a partially collapsed lung, and I was in the middle of nowhere without any real rescue assets—it was then that I realized I was going to die out there.' (Hazlett 2010: 501)

By the linguistic methodology, it should follow from those linguistic data that propositional knowledge can be non-factive, to wit, knowledge does not have to entail truth. This corollary of the linguistic methodology directly contradict the standard understanding of propositional knowledge, that is, truth is necessary for knowledge. This traditional view that knowledge has to be factive reflects in almost all, past or present, mainstream accounts of knowledge (e.g., JTB theory, truth-tracking theory, relevant alternative theory, anti-luck epistemology, virtue epistemology). Hence Harrison claims that:

'The fact that one cannot say such things as 'I know, but I might be mistaken' shows that when one is claiming to know that some proposition is true, one is claiming that there is no possibility of one's being mistaken about its truth' (Harrison 1978: 137)

Therefore, an irreconcilable contradiction between the linguistic methodology and the standard truth condition of knowledge will emerge if Hazlett is right in that Factivity is falsified by non-factive uses of 'knows' that he listed. It is noteworthy that, by listing those non-factive examples, Hazlett does not purport to discard the standard truth condition of knowledge. Instead, what he challenges is just the widely granted linguistic method of analysing epistemological concepts, which he characterises as follows:

‘[T]he linguistic method: imaginary situations are described, and intuitions (which are presumed correct) are elicited concerning whether or not a character in the story said something acceptable’ (Hazlett 2010: 497-98).

Hazlett’s suggestion is not to abandon the standard truth condition of knowledge or to deny those non-factive uses of ‘know-that’ constructions, but a divorce of the linguistic method and the theoretical nature of epistemological concepts:

‘What I’m claiming is that epistemologists have every right to insist that knowledge (as they understand it) is factive—but the price to pay for this (which many will be happy to incur) is to give up the linguistic method described above. I’m suggesting, in other words, a divorce for the linguistic theory of knowledge attributions and traditional epistemology.’ (2010: 500)

Cross-language evidence can also be employed to cast doubts upon the linguistic methodology. Bac & Irmak (2011) point out that phrases like ‘knows wrongly’ (‘yanlış bilmek’) are commonly used among speakers of the Turkish language. The Turkish language accepts statements like ‘He knows that *p* wrongly’ as natural utterances. Nonetheless, this does not mean that Turkish epistemologists have to discard the standard factive conception of knowledge. The linguistic fact that ‘know that’ is used in a non-factive manner does not entail that knowledge-that *per se* is a non-factive concept. There is a gap between an epistemic concept’s linguistic uses and its conceptual nature.

Hazlett does not deny that, at least typically, we use ‘knows’ to guarantee the truth of some propositions. He argues that this only indicates that ‘S knows that *p*’ typically *implies* that *p* is true, rather than *entails* that *p* is true—and this implication can be accommodated by a Gricean explanation. Drawing on Grice (1989), Hazlett argues that knowledge ascriptions such as ‘S knows

that *p*' typically implies *p*'s truth because it is generally mutually assumed by speakers that they conform to at least three maxims of conversational cooperation:

Quality: 'Do not say anything you believe to be false, or which you don't have reason to believe is true';

Quantity: 'Make your contribution to a conversation as informative, and only as informative, as is required';

Relation: 'Make your contribution to a conversation relevant' (see Hazlett 2010: 511-512).

Hazlett argues that the phenomenon that 'knows' typically implies truth can be explained by the fact that we ordinarily obey these Gricean maxims, and thus would only utter something that we believe to be true when making knowledge ascriptions. For example, when asked 'Any information from the FBI about how the bomb was constructed?', if one answers that 'They know the bomb was homemade', now that we assume the speaker is conforming to Quality, the speaker is assumed to be asserting something that she believes to be true. Moreover, given Quantity and Relation, the speaker is also assumed to wish her interlocutor to believe that it is true that the bomb was homemade—otherwise, the speaker should say something like 'They thought they knew that the bomb was homemade, but that was false'. In summary, the speaker's utterance implies that 'it is true that the bomb was homemade'.

Hazlett's non-factive theory of 'knows' invites many objections. Before discussing some most representative objections in more detail, I want to emphasise two things. First, almost all critics of Hazlett admit that his examples involving non-factive uses of 'knows' are authentic rather than invented (in fact, they are all quotations from real publications). What they question is just how to explain those authentic linguistic data. Hazlett argues that there is a simple, effective, and elegant explanation—to wit, 'knows' does not entail truth. Hazlett's opponents reject this non-factive account of 'knows' and argue that those Hazlett-friendly data can be explained away without sacrificing Factivity. Second, the fact that the linguistic nature of 'knows' is controversial (as the debate around Hazlett's proposal has shown)

is sufficient to undermine the linguistic method that absolutists rely on. It is in dispute whether ‘knows’ (linguistically speaking) is factive, however, in contrast, it is much less controversial that truth is necessary for constituting *knowledge* (epistemologically speaking). This asymmetry in controversiality can already provide us with a ground for doubt about the linguistic method. For example, many critics of Hazlett (e.g., Turri 2011a; Tsohatzidis 2012; Dahlman 2016; Domaneschi & Di Paola 2019) attempt to account for the data that Hazlett’s cites by appealing to a polysemous reading of ‘knows’ to the effect that ‘knows’ has two different senses—a factive sense and a non-factive sense. This polysemous attitude towards the semantics of ‘knows’ is common in this debate. However, it would be much more surprising if philosophers also concede that the conceptual nature of knowledge comes in two distinct senses: for one sense the truth condition obtains, while for another sense truth is not necessary for constituting knowledge. Quite the contrary, just as Laurence Bonjour points out, the truth condition is “something almost no philosopher has seriously disputed” (2002: 32). Furthermore, almost all opponents of Hazlett endorse a univocal attitude that the truth condition of knowledge should be preserved, at least for the concept of knowledge that epistemologists are interested in (see Turri 2011a: 143; Tsohatzidis 2012: 449; Hannon 2013: 364; Buckwalter 2014: 391; Dahlman 2016: 150; Domaneschi & Di Paola 2019: 102). This shows that the epistemological nature of knowledge is not as faithfully reflected in the linguistic traits of knows as the linguistic methodology presumes.

With these being said, now I will briefly discuss standard objections to Hazlett’s non-factive theory of ‘knows’. My conclusion is: none of those objections can adequately refute Hazlett’s proposed divorce for the linguistic method and the epistemological investigation of knowledge. So how do opponents of Hazlett account for recalcitrant data for Factivity that Hazlett cites? One standard approach is to resort to Richard Holton’s theory of *protagonist projection* (see Tsohatzidis 2012; Hannon 2013; Dahlman 2016; Domaneschi & Di Paola 2019). Holton (1997) writes that:

'I suggest that these sentences work by projecting us into the point of view of the protagonist; let us call the phenomenon protagonist projection. In each case the point of view into which we are projected involves a false belief. We describe the false belief using words that the protagonists might use themselves, words that embody their mistake. So we deliberately use words in ways that do not fit the case.'
(1997: 626)

In accordance with the protagonist projection theory, when making knowledge ascriptions non-factively, we are just projecting us into the points of view of the protagonists who thought they knew that p where p is in fact false³¹. A central idea of protagonist projection is that we *deliberately* use 'knows' falsely without literally meaning that the protagonist really knows a false proposition. Accordingly, we can explain away Hazlett's data without sacrificing Factivity or admitting that those knowledge ascriptions involving false beliefs are literally true.

Hazlett objects the protagonist projection theory by doubting whether people who assert sentences that he cites are 'deliberately' uttering something that they know to be false—it is unclear that people 'know' that nothing false can be known. Hence it is unclear that speakers who use sentences like 'everybody knew that stress caused ulcers' believe that this knowledge ascription is false. Hannon (2013) casts doubt upon Hazlett's judgment by

³¹ The protagonist projection argument is also claimed to receive empirical support according to Buckwalter (2014) and Domaneschi & Di Paola (2018). Both works report that more participants find the more natural interpretation of utterances involving non-factive uses of 'knows' is not that 'the subject really knows that p ' when p is untrue. Instead, more participants find it more natural to interpret those utterances as 'the subject simply believe that p ' (Domaneschi & Di Paola 2018), or 'the subjects thought they knew' (Buckwalter 2014). The problem of these experiments, as Domaneschi and Di Paola admit, is that they fail to show that 'they simply believe that p ' or 'they thought they knew that p ' entails 'they do not really know that p '. Thus the two options of their surveys are not mutually exclusive. Given Hazlett's Gricean account, it would also be unsurprising that people see 'simply believe' or 'they thought they knew' as the more accurate expressions. But this cannot prove that people deny the *literal truth* of those knowledge ascriptions involving false beliefs. Thus, if those experiments are to provide solid support for Factivity, it must be assumed that 'simply believe' and 'thought they knew' *are not* 'really know' -- but this seems to be just presuming that Factivity is true. Otherwise, for what reason can one conclude that those protagonists just 'do not really know' but 'simply believe'?

arguing that it seems to be very natural for speakers to reply negatively when asked 'Did people really know that stress caused ulcers?' However, this does not undermine Hazlett's claim that it is not obvious that when people use 'knows' non-factively, they are *deliberately* using 'knows' in a way that they believe to be false. Here, we can have two readings of 'deliberately use words in ways that do not fit the case'. According to the strong reading, it requires the speakers to *know* that they are using 'knows' in a false way (this seems to be the reading that Hazlett adopted), to wit, they *know that* Factivity is true. If that is the case, then there is one further defence that Hazlett could have made (though it seems that he did not): champions of Factivity would be begging the question if they presume that Factivity is *known* to be true. As for the weak reading of the deliberateness, it also requires people to *presume* that Factivity is true, but it is unclear whether this presumption is commonly shared by ordinary people.

The second line of objection that Hazlett takes is to appeal to the principle of charity. It is argued that if the explanation provided by champions of protagonist projection theory is true, then they would impose systematic falsehood on our ordinary talks of 'knows'—every non-factive use of 'knows' would thus be subject to this falsehood. In contrast, if 'knows' is interpreted as not requiring factivity, and we also endorse that when asserting 'S knows that *p*' we do typically imply that *p* is true (but this implication does not *entail* Factivity, according to Hazlett's Gricean explanation), then the systematic falsehood will not be found in the non-factive theory of knows. Thus, a non-factive account of 'knows' enjoys the advantage of charity over the traditional factive one³².

Another standard objection to Hazlett appeals to the idea that 'knows' is polysemous. Hazlett admits that he cannot prove this idea to be false; he just

³² Most opponents of Hazlett do not deny this. Hannon (2013) only remarks that this extra systematic falsehood is not acceptable if we have justifiable reasons to impose it. Thus, Hazlett's non-factive theory at least enjoys a *ceteris paribus* advantage over its opponents on this aspect.

insists that a polysemous account of knows is inelegant and cannot make the traditional factive definition of knowledge *the* definition, other than one among many. In addition, as we have noted before, the claim that 'knows' has two distinct senses already suffices to cast doubts upon the linguistic method as the epistemologically interesting concept of knowledge is universally accepted as univocal: it must entail truth. On this point, Tsohatzidis (2012) suggests that *factive knowledge* is the chosen subject matter of epistemology, while the fact that knows has another non-factive sense is linguistically irrelevant to this choice. Hazlett (2012) objects that, if that is the case, then the fact that knows has a factive sense should also be linguistically irrelevant to the traditional epistemologist's choice of subject matter. In other words, this simply undermines the linguistic method that the linguistic traits of 'knows' can give our judgments of knowledge's epistemological characteristics a solid ground (see also Hazlett 2010: 519).

Now we have seen how we can divorce the linguistic uses of 'know that' from the *standard* epistemological conception (e.g., knowledge has to be factive). In a similar vein, this divorce can also be accomplished by considering some *non-standard* epistemological theories. The first example that I want to discuss is the non-doxastic account of knowledge. I do not think this case study could be as intuitive as the non-factive case discussed above for all my readers, as the non-doxastic account of knowledge is somewhat counterintuitive (but couldn't the minority turn out to be correct?). However, I believe that it could still help to illustrate how the divorce for linguistic methodology and *minority* epistemology can be realised. The traditional view holds that belief is a necessary for propositional knowledge, 'knowing that *p*' entails 'believing that *p*'. On the contrary, Radford (1966) argues that knowledge can be independent of belief by constructing a scenario where an unconfident student correctly answers many questions asking her to write down the date of a certain historical event (e.g. Queen Elizabeth died in 1603) in her history exam, though she does not actually believe in her answer as she feels like she was just guessing. Radford argues that although

the unconfident student lacks relevant belief, it is still acceptable to claim that she knows that Queen Elizabeth died in 1603 as her answer is correct and epistemically more valuable than a mere guess. Many epistemologists do not agree with Radford that the unconfident student does possess knowledge, so Myers-Schulz & Schwitzgebel (2013) modifies Radford's case. In the modified unconfident student case, the student is described as a well-prepared examinee who had a good revision before the exam. While in the last minute of the exam, she got nervous and unconfident as she realised that the time was running out. Therefore she wrote that 'Queen Elizabeth died in 1603' in disappointment and self-doubt, which is exactly the correct answer. Basing on this modified case, Myers-Schulz & Schwitzgebel conducted a survey among undergraduate students asking whether the protagonist has belief and knowledge. A divorce for 'believes' and 'knows' can thus be found in terms of people's epistemological intuition.

Apart from that, Farkas (2015) also argues that belief might not be a necessary condition for knowledge by trading on the extended mind thesis. Farkas claims that there are some convincing cases where propositional knowledge can be extended by lights of the 'extended knowledge' proposal, *e.g.*, the famous Otto's case put forward by Clark & Chalmers (1998). In that envisaged scenario, Otto can be taken to know that 'MoMA is on 54th street' (call the proposition 'M') in virtue of the information recorded in his coupled notebook. Nevertheless, Farkas argues that there are sufficient reasons to deny that Otto has the belief that M. For example, one reason that she refers to is from Adams & Aizawa (2010)—mental states (such as beliefs) need to have non-derived intentionality, which is what Otto's record of M in his notebook lacks. In light of this, Farkas concludes that Otto's case can be recognised as a case of knowing without believing.

If those arguments for non-doxastic epistemology are right, then there would be a contradiction between the (belief-unnecessary) nature of knowledge and our daily uses of 'know-that' in the ordinary language. That is because, no

matter the non-doxastic epistemology is granted or not, utterances as follows are still intuitively inconsistent:

(17)??? I know that Queen Elizabeth died in 1603,
but I do not believe that she died in 1603.

(18)??? I know that MoMA is on 54th street, but I
do not believe that it is on 54th street.

Let me reclaim that I do not aim to defend the non-doxastic epistemology and conclude that belief is unnecessary for knowledge. I am just showing how our linguistic evidence could contradict the epistemological nature of the concept 'knowledge', *if the non-doxastic epistemology is right*.

Likewise, drawing on Wittgenstein's reading of *hinge commitments*, some hinge epistemologists such as Moyal-Sharrock (2004), McGinn (1989) and Pritchard (2012b)³³ have argued that our hinge commitments to basic quotidian propositions (e.g., 'I have hands', 'The Earth has been in existence for a long time before I was born') cannot constitute knowledge-apt *beliefs*—we do not virtually hold rational *beliefs* of those hinge propositions. That is because beliefs, qua a type of propositional attitude, are responsive to rational considerations, but hinge commitments are visceral and 'animal' in nature—they are more like ways of acting rather than beliefs. Hence Pritchard (2012b) argues that:

'But that is just to say that it is not the kind of commitment that could ever be responsive to rational considerations, whether in its favour or against it, and it is hard to see how a commitment of that sort could be properly characterised as a belief. (A belief, after all, is a belief that such-and-

³³ Pritchard's non-epistemic reading differs from the non-propositional reading of Moyal-Sharrock and McGinn in the sense that the non-propositional reading does not endorse that hinge commitments are propositional attitudes – while the non-epistemic reading endorses so, but just insists that they are not beliefs. Nonetheless, what is common for both the non-epistemic reading and non-propositional reading is that they refuse to recognise hinge commitments as beliefs.

such is so, and hence is in its nature in principle responsive to rational considerations—i.e., considerations which indicate whether something is so. That is why wishful thinking, even when in all outward respects very alike belief, is not belief’ (2012b: 267-68).

One somewhat surprising upshot of this Wittgensteinian view is that we do not indeed *know* those hinge propositions now that we do not have beliefs (but just groundless hinge commitments) of them. Hinge propositions are not in the market for knowledge as they cannot be virtually *believed* (here, unlike the non-doxastic epistemology discussed above, the Wittgensteinian view still sees belief as a necessary condition for knowledge). Hence, presume that this Wittgensteinian view is right, then our linguistic evidence would be standing on the opposite side, as it is felicitous and natural regardless to say that ‘I know I both hands’ or ‘I know that the Earth has been in existence for a long time’ in our ordinary language³⁴. If one sticks to his guns and insists on the linguistic method, then the Wittgensteinian view is unacceptable anyhow as it predicts something significant incompatible with our linguistic data and linguistic intuition.

The linguistic argument for absolutism postulates the linguistic methodology that how we use a notion in the ordinary language determines how we should understand the notion’s conceptual traits. This methodology appears to be intuitively legitimate as we ordinarily talk of ‘gradability’ from a linguistic point of view. Notwithstanding, I have explained why we had better discard the linguistic methodology in terms of both standard and non-standard

³⁴ A contextualist reading of the Wittgensteinian hinge epistemology grants that there are daily contexts where it is natural to assert those claims, and hence those ‘Moorean propositions’ might potentially be felicitously claimed to be known. For example, contexts of daily communication. But these will not be contexts where those claims functioning as ‘hinges’. According to contextualism, those claims typically function as hinges in contexts where they are called into question by sceptics. In those contexts, these claims cannot be known or rationally doubted (see Pritchard 2011). However, this reading does not affect my argument. That is because, even in contexts where scepticism is involved, it is still (at least linguistically) felicitous to claim that ‘scepticism is wrong, as I know that I have both hands’. Wittgensteinian contextualists might find those claims inappropriate as they violate an essential conceptual property of hinges (*i.e.*, the unknowability), but this inappropriateness (provided that the Wittgensteinian view is correct) is philosophical, rather than linguistic.

epistemological understandings of knowledge. It has been revealed that there is an explanatory gap between linguistic evidence and the epistemological nature of propositional knowledge that absolutists can hardly cross. Why does this gap exist? I think there are at least three possible explanations. First, as we have seen before, the linguistic methodology turns essentially on people's linguistic intuitions, while intuitions regarding how 'knows' should be properly used might vary from people to people, or from case to case. In contrast, it is a natural philosophical aspiration that 'knowledge', *qua* the central concept of epistemology, ought to be of a more stable conceptual essence. Second, it is evident that different languages can have fundamentally different 'conventional usages' of 'knows'. However, it seems to be much less evident and persuasive that *the notion of knowledge* also differs so fundamentally for speakers of different languages. Insofar as we see knowledge as an elemental notion for all human beings, it would be also natural to expect that some basic epistemological consensuses (e.g., knowledge entails truth) can apply in a cross-language manner rather than be exclusive to English-speaking countries³⁵. Last but not least, the gradability of knowledge can be evaluated in more dimensions other than the mere linguistic one. For example, a more philosophically interesting way of discussing knowledge's gradability should concern whether knowledge can be improved and how the quality of knowledge can be evaluated. Hetherington (2011) also expresses qualms about epistemologists' worship of the linguistic methodology:

'Much traditional epistemology is said to involve conceptual analysis — uncovering meanings, perhaps ascertaining essences, and so deriving some significant truths. Often, the methodology that is used apparently relies upon claims about how a word such as 'knowledge' is generally used. Care is needed, of course, because epistemologists, while confidently relying on various claims as to how people use a given term, might dismiss competing uses of those terms or associated ones as being confused, for example. Is there

³⁵ Admittedly, absolutists can be steadfast and assert that speakers of different languages have different conceptions of knowledge – but this would also undermine the applicability of the absolutist conclusion.

ever a danger of epistemologists bolstering their arguments by respecting only those linguistic usages that cohere with their own ways of speaking? Surely there is.' (2011: 149)

With these being said, the linguistic methodology should not be adopted so unhesitatingly. Instead, discarding the worship of it can lead us to a more fruitful discussion regarding the gradability problem of knowledge.

In summary, this section provides three reasons why the argument from linguistic evidence (which plays an essential role in people's accepting *overall-absolutist attitude* or the *equivocal-absolutist attitude*) cannot justify epistemic absolutism: 1) there are competing linguistic data for epistemic gradualism in the English language; 2) there are also conflicting linguistic intuitions in other languages; 3) *even if* the linguistic evidence favours absolutism over gradualism in general, the linguistic methodology that absolutists adopt does not suffice to reflect the conceptual nature of knowledge-that. Now let us move on to address another important argument for absolutism.

2. Objections to the Argument from Object

Knowledge has its object. The object of an item of knowledge-who is a person; the object of an item of knowledge-where is a location. But what is the object of propositional knowledge? A straightforward and simple answer is: the proposition p is known. According to the *well-known* T-schema³⁶ (see Tarski 1933, 1944, 1983; Dummett 1959), the proposition p is true if and only if p . Therefore, to say that 'snow is white' is equivalent to say that "snow is white' is true'. Being based on the T-schema, the standard view holds that 'know that p ' is equivalent to 'know that p is true', which means that the object of knowledge-that- p is the truth of the proposition p . This understanding of knowledge might be justifiable if one endorses the T-schema. There seems to be a natural derivation as follows:

³⁶ Also known as 'the equivalence schema'.

[T-schema] $\langle p \rangle$ is true iff p .

(Where 'iff' expresses the material biconditional, and the angle brackets indicate an appropriate name-forming device meaning 'the proposition that...'. For the sake of simplicity, I will omit angle brackets when formulating KT-schema.) Correspondingly, one might derive:

[KT-schema] S knows that p , iff, S knows that p is true.

Given that ' p ' is equivalent to ' p is true', it would make *prima facie* perfect sense to substitute ' p is true' for ' p ' in the know-that construction, hence KT-schema is intuitively appealing. Another way to understand why KT-schema is intuitive is to resort to the closure principle. The closure principle predicts that if S knows that p , and that p entails that q , then S knows that q . A proof of KT-schema can thus be constructed:

- (i) S knows that p .
- (ii) S knows that ' $p \leftrightarrow "p" \text{ is true}$ '.
- (iii) The closure principle.
- (iv) S knows that $p \leftrightarrow$ S knows that p is true.

Few epistemologists have systematically argued for KT-schema, perhaps because this schema is too intuitive. Nonetheless, the traditional understanding of propositional knowledge behind KT-schema was reflected in epistemologists' interchangeable uses of 'know that p ' and 'know that p is true'. For example, Dretske (2014) defines the closure principle as follows:

'Closure is the epistemological principle that if S knows that P is true and knows that P implies Q, then, evidentially speaking, this is enough for S to *know that Q is true* [emphasis added]' (Dretske 2014: 27).

Then in the next page of the same paper, Dretske formulates the closure principle as:

'S knows P.
S knows P implies Q.
Therefore, S *knows* Q [emphasis added].'
(Dretske 2014: 28)

Here, 'knows Q' and 'knows that Q is true' are used interchangeably. Another evidence is from Ram Neta:

'What's the point of trying to give an account of the necessary and sufficient conditions for a subject S *to know the truth of some proposition p* (what I'll henceforth call an 'analysis' of '*S knows that p*') [emphasis added]?' (Neta 2002: 663)

It is also not hard to see how 'knows the truth of some proposition *p*' and 'knows that *p*' are deemed as identical here. I cannot exhaust all instances here³⁷, but it is noteworthy that this understanding of propositional knowledge is also embodied in some influential epistemological theories. Truth-tracking theory of knowledge (see Nozick 1981; Dretske 1971; Goldman 1967; *etc.*) claims that to know that *p* is to track the truth of *p*. With the KT-schema in play, it will be easy to comprehend tracking theorists' emphasis on the truth, as tracking the truth is tracking the object of knowledge.

Hetherington (2005) argues the 'strongest form of defence' for epistemic absolutism, *i.e.*, the argument from object, is based on the idea that to know that *p* is to know that proposition *p* is true. If knowing a proposition is knowing the truth of that proposition, then propositional knowledge has to be absolute in the sense that the truth of a proposition cannot be known more or less—it can only be known as an absolute whole. Epistemic absolutism can thus be defended by virtue of KT-schema. Hetherington does not elaborate on the

³⁷ For more instances, see Dretske 1971; Moser 1987:91; Moffett 2003:82; Hetherington 2005:148; Jespersen 2008:125; *etc.*

argument very clearly, however, we can try to formulate the argument as follows:

[Argument from Object]

(AO1) To know that p is to know that p is true.

(AO2) The truth of a proposition p is absolute as it cannot be known more or less.

(AO3) The truth of a proposition p is absolute as it cannot be known better or worse without thereby altering the object of the original knowledge that p .

(AOC) Propositional knowledge is absolute.

Hetherington's own objection to this argument from the object is rejecting KT-schema, viz, arguing that 'knowing that p ' is *more than* 'knowing that p is true'. However, his objection relies heavily on his controversial radical anti-intellectualism, which claims that knowledge-that is a kind of how-knowledge—knowing how it is the case that p . On the basis of his how-knowledge theory, Hetherington argues that:

'[K]nowledge that p is almost always knowledge of more besides. It generally includes knowledge of further aspects of how it is that p —aspects beyond p as such, most narrowly construed. On such occasions, to know that p is to know more than p , to some extent or other. It is to know truths, for instance, other than p —even as part of knowing that p . For it is to know some more or less extensive array of aspects of how it is that p .'
(Hetherington 2005: 148-49)

Hetherington's how-knowledge theory will be introduced in more detail in the next section, and it will be revealed that this theory is severely problematic. Apart from that, Hetherington's objection to the argument from the object is not satisfactory in the sense that he is simply asserting that knowing that p is knowing how it is the case that p rather than knowing that p is true. In other words, at most he just proposes an alternative account of the object of

propositional knowledge that competes with the traditional account supported by KT-schema, rather than shows us why KT-schema is false. He does not virtually disprove KT-schema—no substantial attack on KT-schema has been given. Here, I aim to provide such a substantial attack by presenting counterexamples to KT-schema in which one knows that a proposition p is true, but does not know that p .

2.1: Objection to the KT-Schema³⁸

The objection that I will offer here captures the sense that knowing that p also requires understanding the semantic meaning of p (to a certain extent), while a subject who is (completely) ignorant of the semantic meaning of a proposition p can still know that p is true. In other words, one can know a proposition that he completely does not understand to be true, while one cannot know a proposition that he completely does not understand. A most straightforward example is, one can know a proposition written in a foreign language to be true (on the basis of testimonies or logical reasoning, *etc.*), while still fails to know the fact expressed by the proposition. Consider the following case:

Math Textbook Suppose that your math teacher told you that 'Every proposition in your textbook is true'. Your teacher is reliable. Your textbook is classic and well-proofread. You believe that testimony firmly as there is no proper reason to doubt it. Given these, your belief constitutes knowledge in that case, *viz*, you know that every proposition in your textbook is true. As a result, you also know that the last proposition of your textbook is true. However, you have not read the last

³⁸ Elsewhere, I have provided a more detailed objection to the KT-schema which includes a more technical dilemma for the KT-schema. See Lai (2019).

proposition of the textbook, so what that proposition expresses is completely unknown to you.

That is a case where you know the truth of a proposition without knowing that content of the proposition. In the math textbook case, one fails to know the last proposition of the textbook because the content of the proposition *per se* has not been seen or be accessed in any way by the subject. Apart from cases like this, there are also cases where the content of the proposition is seen. For example:

Foreign Proposition There is a true Chinese sentence ‘爱丁堡是苏格兰首都’ (which means ‘Edinburgh is the capital of Scotland’) printed in a pub quiz book. A non-Chinese speaker, Jonny, read this Chinese sentence, but he did not understand that sentence at all. So he asked his Chinese friend Chen, a reliable man who lives in Scotland, for help. Chen told Jonny that this sentence is true. Jonny believed Chen’s answer. The next question of the book is ‘What is the capital of Scotland?’ Jonny spoke to Chen: ‘Yeah, I know this one. The answer is Glasgow!’

In this case above, I take it to be clear that Jonny can be granted as knowing the proposition ‘爱丁堡是苏格兰首都’ is true by virtue of Chen’s testimony.

This should be a typical case in which one’s true belief justified by testimonies constitutes knowledge. However, it is obvious that Jonny does not know that Edinburgh, rather than Glasgow, is the capital of Scotland.

Therefore, it will be highly implausible to admit that he virtually knows that 爱丁堡是苏格兰首都. KT-schema is thus violated. The example above involves hybrid languages expressions, which might look not straightforward enough. Consider a modified case:

Foreign Proposition 2 There is a true English sentence ‘Edinburgh is the capital of Scotland’ printed in a pub quiz book. A non-English speaker, Chen, read this English sentence, but he did not understand that sentence at all—he does not understand any single word of that sentence. So he asked his Scottish friend Jonny, a reliable man who can also speak Chinese, for help. Jonny told Chen that this sentence is true. Chen believed Jonny’s answer, but he still did not understand that English sentence. In fact, Chen just does not know what the capital of Scotland is—even if the question was asked in his mother language. He just does not know there is a city in the world named Edinburgh (or ‘爱丁堡’ in Chinese).

In this modified case, Chen can also be recognised as knowing that the proposition ‘Edinburgh is the capital of Scotland’ is true by virtue of Jonny’s testimony. Readers who grant that Jonny knows the proposition at issue to be true in our first foreign proposition, should also be able to endorse that Chen’s belief can constitute knowledge in the same manner. Whilst knowing that the proposition ‘Edinburgh is the capital of Scotland’ is true, Chen does not know that Edinburgh is the capital of Scotland. The content of the proposition is not transparent to him. He is completely ignorant of the relevant fact that the proposition supposes to deliver. It is almost incoherent to say ‘I know that Edinburgh is the capital of Scotland, but I do not know what the capital of Scotland is’³⁹. Chen does not know what the capital of Scotland is. He simply does not know that there is a city named Edinburgh, which is the capital city of Scotland—no matter the fact is expressed in English or Chinese. KT-schema is thus disproved in that case.

³⁹ In contrast, it is coherent to say that ‘Chen told me that ‘爱丁堡是苏格兰首都’, I know that must be a true statement, but I do not understand that sentence, and I am still ignorant of what the capital of Scotland is.’

In both foreign proposition cases, the proposition in question (‘爱丁堡是苏格兰首都’ and ‘Edinburgh is the capital of Scotland’) is just like a string of mysterious codes for our protagonist. The semantic content of the string of codes is not important, because it does not influence the protagonist’s judgment of the proposition’s truth value. Our protagonists know those propositions to be true, just on the basis of a true reliable testimony, rather than the semantic content of those propositions *per se*. They only have knowledge *de dicto*, rather than knowledge *de re*.

We have seen that, if the subject does not understand the meaning of words constituting the proposition in question, then it is possible that KT-schema will be violated. Besides, lacking understanding of the proposition’s logical structure might also lead to the failure of KT-schema. Consider another example:

Logic Exam There is a multiple choice question in a logic exam asking candidates to choose the only one correct answer from three options. A student knew clearly that both option A and option B are false, while he did not understand what does option C ‘ $p \wedge q$ ’ mean. That is because the student did not remember whether the symbol ‘ \wedge ’ represents ‘and’ or ‘or’. Nevertheless, the student still decided to choose C confidently, as he knew that A and B cannot be the correct answer, and thus option C has to be true. In fact, C is the correct answer indeed. However, without knowing that ‘ \wedge ’ is the conjunction sign, when asked ‘whether p or not’ and ‘whether q or not’ after the exam, the student answered ‘I don’t know’.

In this case, I believe that most of us could admit that the student knows that ' $p \wedge q$ ' is true. The student came to know this by the method of exclusion, *i.e.*, he inferred that option C is true from the premises that: (i) neither A nor B is true; and (ii) one of the three options has to be true (presume that there is no sign indicating that the test question is misprinted or provides three wrong options). His reasoning is valid and sound, therefore it should be safe to conclude that the student knows that ' $p \wedge q$ ' is true.

On the contrary, it will be much more improper to claim that the student knows that $p \wedge q$. That is because the student does not really understand the logical structure of the proposition in question, and thus fails to comprehend the semantic meaning of the proposition. Moreover, it is highly likely that the student would choose option C regardless, even if the proposition expressed by option C were ' $\sim p \vee \sim q$ '. The content of option C is not important for the student to choose C. For that student, option C is just a 'place-holder' rather than a virtually meaningful proposition whose semantic meaning would affect its truth value. No matter what proposition appears in the place of option C, the student would deem it as true anyway. In addition, it will be extremely counterintuitive to claim that one can know a conjunctive proposition without knowing any single conjunction. It is incoherent to claim that 'I know that James is a boy and Jimmy is a boy, but I do not know that James is a boy'. However, in the logic exam case, the student does not know either p or q, which means that he does not know any conjunction of the conjunctive proposition. In contrast, the following statement is coherent: 'I know that the proposition ' $p \wedge q$ ' is true, but I do not understand the meaning of the proposition, therefore I do not know whether p is true or not'. The upshot is, we can hardly grant that the student knows the conjunctive proposition, even though he knows (*de dicto*) that the proposition is true. Therefore, KT-schema is violated again in this case.

Let us recap all those counterexamples briefly. The math textbook case concerns a situation where the content of the proposition is *completely*

unrevealed to the subject and thus KT-schema is violated. The foreign language case involves a scenario where the content of the proposition is *exposed* to the subject, nevertheless its semantic meaning is *completely mysterious* to the subject and hence KT-schema fails again. The logic exam case further constructs a scenario where the content of the proposition, whose semantic meaning is *partially understood*, is *exposed* to the subject. However, because of the incomprehension of the key logical structure of the proposition, the subject still fails to know the proposition and thereby falsifies KT-schema. The extent to which the protagonist of the three types of counterexample understands the content of the given proposition seemingly increase by degrees. Nevertheless, what those protagonists have in common is that their understandings are not enough. Admittedly it will be an interesting question to ask how much understanding is enough, but that is not the main problem that this chapter aims to answer. The moral that I hope to draw from those counterexamples is: (a certain extent of) understanding of the content of a proposition p is necessary for knowing that p , while it is not necessary for knowing that p is true. With this distinction in play, we can better understand why there is a gap between [knowing that p is true] and [knowing that p], and why KT-schema fails. So, if knowing the truth of p is insufficient for answering the object problem of propositional knowledge, what are we claiming to know when we claim to possess an item of knowledge-that- p ? From what we have discussed, a potential answer might be that: 'to know that p ' seemingly means that 'to know that p is true + to know the content of p '. To be specific, we reject:

[T Account] To know that p is to know that p is true.

But advocate:

[T+C Account] To know that p is to know that p is *true*, plus, to know p 's *content*, which requires understanding the meaning of p 's content.

There are three points that I wish to clarify regarding this T+C account. First, this account purports to answer the object problem rather than the definition problem of propositional knowledge. The object problem *differs from*, albeit closely-related to, the definition problem. Theories attempting to answer the definition problem include the historical 'JTB' template, the relatively recent 'JTB + anti-Gettier factors' template, and the AAA-model of virtue epistemology, *etc.* However, none of them is answering the object problem of propositional knowledge. Any attempt to solve the object problem should be careful about the subtle distinction between the two problems. Given this distinction, one should not worry that our 'truth and content' account will invite circular definition or infinite regress. That is because, we are not *defining* 'knowing-that' by 'knowing the truth and the content of *p*'. We are just *analysing the object* of 'knowing-that' by the 'truth and content' account.

Second, the T+C account parallels 'knowing the truth of *p*' with 'knowing the content of *p*' as two factors constituting the object of knowledge-that-*p*. This does not mean that these two factors are completely independent of each other. They could be interdependent and even overlapped to some extent. We have seen how one can know the truth of *p* without knowing the content of *p*. It is also easy to imagine scenarios where one knows the content of *p* without knowing its truth. For example, one can comprehend a proposition without believing that proposition. The union of the two factors can constitute a more complete and thus better answer to the object problem, which can at least avoid counterexamples that we discussed before. However, is this T+C alternative the *full* answer to the object problem of propositional knowledge? Is there any other factor constituting the object of propositional knowledge except for 'truth and content'? Or, it may be questioned that, is 'the content of *p*' a necessary component of the object of knowledge-that-*p*? Furthermore, is 'the truth of *p*' really necessary as it is orthodoxly taken to be? I will leave these questions open to my readers. By this stage, my purpose is only to

debunk the T-account and propose a better alternative—it does not need to be the best answer, but it can serve as a promising replacement.

Third, according to Hetherington (2001), there is a close connection between ‘knowing’ and ‘understanding’, which coincides with our conclusion that knowing that p requires understanding the content of p . One can understand the content of a proposition better than another person does. If ‘understanding-that’ partially constitutes ‘knowing-that’, then it is reasonable to infer that knowing-that can also be gradable as understanding-that. That is, it will be defensible to argue that knowledge-that- p can be graded in terms of one’s understanding of the semantic meaning of p —the better the meaning is understood, the better the proposition is known. Admittedly, one cannot derive directly that knowledge-that is gradable from that its *necessary conditions* are gradable. Beliefs and justifications are also ordinarily taken to be necessary for knowledge, nevertheless, their gradability does not imply the gradability of knowledge-that *per se*. Thus the discussion regarding the gradability of understanding should better be understood as an *internal criticism* of epistemic absolutism, which reveals why absolutism does not follow from premises put forward by absolutists. That is, if the ungradability of truth (*qua* the alleged object of knowledge-that) can be employed to defend epistemic absolutism, then the gradability of understanding should be able to undermine absolutism and support gradualism. That is because, we have noted that the semantic content of p , as well as the truth of p , also constitutes the object of knowledge-that- p . I take it to be justifiable that there is no substantial difference between ‘knowing the semantic content of p ’ and ‘understanding the semantic content of p ’. Admittedly, this claim needs to be borne out by more intensive analyses. But if this claim turns out to be true, and insofar as we grant the premise that the gradability of knowledge is influenced by the gradability of knowledge’s object, then given that understanding is gradable, it is also reasonable to argue that knowing that p can also be deemed gradable. This inference does not resort to the gradability of the *necessary condition* of knowledge, but the gradability of the

object of knowledge. Notice that this inference is premised on that we grant the underpinning presumption of gradualists' argument from object, to wit, knowledge's object determines knowledge's gradability. Section 2.2 will provide a different objection that challenges this presumption.

It is hard indeed to determine how much understanding is *enough* for constituting knowledge. Actually, I think it is almost unlikely to locate a non-arbitrary and clear cut-off point between 'sufficient' and 'insufficient' understanding (for constituting knowledge). It is as hard as to locate a non-arbitrary and clear cut-off point for 'enough justification (for constituting knowledge)'. Notwithstanding, this sort of 'boundary problem' should not hinder us from taking gradable terms like 'understanding' into consideration when answering the object problem. After all, an intrinsic trait of most gradable terms is the fuzziness of their boundaries. *Belief* requires confidence, but how confident is confident enough for a belief? *Gloom* requires sadness, but how sad is sad enough for one to be recognised as gloomy? Likewise, clear-cut boundaries of many other gradable concepts—such as 'enjoy', 'hatred', 'devoutness', *etc.*—can hardly be determined in a non-arbitrary way. Hence, if knowledge-that is gradable, then we shall not even be perturbed by the boundary problem.

2.2: Objection to Other Premises

The argument from object claims that the object of an item of knowledge-that-*p* is the truth of the proposition *p*. Since knowledge of the truth of *p* cannot be gradable, knowledge-that-*p* cannot be gradable. Objections to KT-schema in the previous subsection challenge the premise that the object of knowledge-that-*p* is just the truth of *p*. Now let me exhibit another approach to block absolutists' argument from object, which can be seen as an *external criticism* of absolutism. The idea is that even if we grant KT-schema and thus concede (AO1), it is still unclear how *the absoluteness of knowledge's object* could entail *the absoluteness of knowledge per se*. In a more general sense,

it is a *non sequitur* to conclude the absoluteness of an epistemic concept from the absoluteness of this concept's object. Take some paradigmatic gradable verbs, such as 'agree' and 'believe', for examples. It can also be claimed that the object of a propositional belief (e.g., belief that Edinburgh is the capital of Scotland) is the truth of the proposition that 'Edinburgh is the capital of Scotland'; similarly, the object of one's agreement of a proposition is the truth of the given proposition (e.g., to agree that Edinburgh is the capital of Scotland is to agree that Edinburgh is the capital of Scotland is true). Nevertheless, believe-that-*p* and agree-that-*p* are still gradable, which means that merely indicating that the object of ψ -that-*p* (where ψ is a verb) is the truth of *p* cannot entail that ψ is ungradable.

The real substantial argument might lie in (AO2) and (AO3). Unlike 'believe' and 'agree', one cannot say 'I completely know that *p*' (in contrast, 'I completely agree that *p*' is seemingly felicitous) or 'I half know that *p*' (in contrast, 'I half believe that *p*' is seemingly felicitous). The truth of a proposition can neither be known partially, nor more or less; on the contrary, it can only be known completely as a whole. However, as we have noted in Chapter One, (AO2) can only at most show that propositional knowledge is quantitatively ungradable. The question that we concern in this gradualism/absolutism debate is whether propositional knowledge is qualitatively gradable. (AO3) involves some insights into this question. According to (AO3), the object of a given item of knowledge-that-*p* will be altered once one tries to 'improve' one's knowledge. New justification would at best lead to a new item of knowledge-that-*q*. For instance, a person came to know that Jonny is the murderer because she witnessed Jonny committing the murder. Gradualists might argue that her knowledge that Jonny is the murderer can be improved if she also heard Jonny pleading guilty in court. A new justification (Jonny's pleading guilty) is gained in this case. However, proponents of the argument from object would suggest that, rather than improved her original knowledge that 'Jonny is the murderer', the subject at best gained a new piece of knowledge that 'Jonny admitted his guilt in

court'—the object of the original item of knowledge has been altered. The fact that Jonny is the murderer has been known as a whole by the subject when she witnessed Jonny's murder. Jonny's admitting his guilt is just another item of fact which becomes the object of another item of propositional knowledge. That is different from knowledge of a person: knowing more facts of a person would not alter the object of an item of knowledge-who—the object is still that person. On the contrary, knowing more facts (justifying another fact) means altering the object of an item of knowledge-that as you are not knowing the same fact expressed by the original proposition.

However, the claim that 'new justifications would alter the object of the original item of knowledge-that- p ' is untenable. Even if we grant that the object of an item of knowledge-that- p is the truth of p , it is still unclear why one cannot know the truth of p better by virtue of more sufficient justification. Absolutism cannot just identify 'knowing that p ' with 'knowing that p is true', and then conclude directly that ' p is true' cannot be known better or worse. That is just begging the question, as ' p is true' is also a proposition. It would also be untoward to insist that new justification would lead to new propositional knowledge. One came to know that Edinburgh is the capital of Scotland from her geography textbook. However, it would be odd to claim that if Wikipedia also told her that Scotland's capital is Edinburgh, then this justification provided by Wikipedia can only provide her with another item of knowledge that 'Wikipedia says the capital of Scotland is Edinburgh', rather than improve her original knowledge that 'the capital of Scotland is Edinburgh'. The fact delivered by the original proposition is still the object that the subject concerns with, and it is still the object that one's new evidence is supposed to justify.

Apart from the fact that absolutists have never provided any independent argument elaborating why *knowing more p -relevant facts* cannot improve one's original knowledge-that- p , another important defect of this explanation

of (AO3) is that there is more than one way in which knowledge can be improved. The above explanation of (AO3) only deals with the kind of epistemic improvement achieved by *gaining more evidence*. However, knowledge can be graded in accordance with other criteria, and epistemic improvement can be realised in many other manners. For instance, reliabilists can argue that one item of knowledge-that-p can be improved if the relevant belief is formed in a more reliable way. In this case, *more evidence* is not necessary, rather, a better way to form the belief is what reliabilism concerns. Similarly, virtue epistemologists can also define the improvement of knowledge under the AAA-model: one can know that *p* better if one's relevant epistemic performance is more accurate, more adroit and apt—degrees of accuracy, adroitness and aptness does not necessarily concern more justification. In summary, (AO2) can only defend the quantitative absolutism, which is not what we aim to object to. (AO3) does not hold water and seems to be a *petitio principii*. The argument from object cannot defend epistemic absolutism successfully even if KT-schema is granted.

3. Objections to the Argument from Contextualism

Responses to this argument for absolutism should not be the focus of this chapter, as this argument, by its nature, is a 'conditional censure' to epistemic gradualism that blames epistemic gradualism be based on contextualism. Surely, one can choose to reject this argument by defending contextualism directly, but that is beyond the purpose of this thesis. Instead, chapter 4 will demonstrate how gradualism can be constructed without resorting to contextualism. If the alleged connection between contextualism and gradualism is cut off, then absolutists cannot easily discard gradualism as an appendage of contextualism. Here I will only provide a rough picture of how epistemic gradualism can be developed without being tied up with contextualism or even any particular account of knowledge, as many influential theories of knowledge can be utilised to construct gradualism:

For the classic JTB template of knowledge, it can be argued that within this account of knowledge, knowledge can be graded in terms of degrees of belief, justification and even accuracy of truth. For proponents of virtue epistemology, gradualists can propose that propositional knowledge can be graded in terms of gradations of epistemic competence and levels to which that epistemic competence is manifested. Sosa's distinction between 'animal knowledge', 'reflective knowledge' and 'knowledge full well' can also be employed to grade the quality of knowledge. In addition, as we have noted before, for externalism such as reliabilism, it can be argued that the quality of an item of knowledge-that- p is determined by the reliability of the process of forming the true belief that p . While when it comes to knowledge-first epistemology, the combination of 'knowledge' and 'evidence' could be explored. It is plausible for knowledge-firsters to argue that the better an instance of knowledge-that- p is, the better it can be an evidence for believing that p , and vice if the quality of 'evidence' is deemed as gradable.

4. Concluding Remarks

This chapter has provided objections to all three mainstream arguments supporting epistemic absolutism. Among them, the argument from linguistic evidence and the argument from object are the focus of this chapter. Both internal and external criticisms have been offered in order to argue against the two arguments. With all three main arguments for absolutism being debunked, it should be fair to conclude that absolutism is not as tenable as the orthodoxy takes it to be. Instead, absolutism is ill-grounded. Therefore, it is time for us to give its opposite, *i.e.*, epistemic gradualism, a more serious consideration. As we have mentioned before, currently, the most developed gradualist account of knowledge is constructed by Stephen Hetherington. The next chapter will critically analyse Hetherington's gradualist theory in detail.

Chapter Three: Hetherington on Epistemic Gradualism

Abstract: In this chapter, I will analyse and criticise Hetherington's two versions of gradualism. The first version is based on his 2001 book *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*, which can be regarded as the most well-known and systematic epistemic gradualist theory that can ever be found in the current literature. The second version of his gradualist proposal is mainly developed in one of his 2005 papers entitled 'Knowing (How It Is) That P: Degrees and Qualities of Knowledge' and his 2011 book 'How to Know'. The first two sections will introduce how Hetherington constructed the two versions of gradualism respectively. Differences between the two forms of gradualism will be highlighted. Section 3 and 4 will call into question Hetherington's two gradualist proposals.

In chapter two, we have seen why arguments underlying epistemic absolutism are all untenable. Now we shall move forward to see how epistemic gradualism can be justified. On our way towards the revitalisation of the long-forgotten debate between gradualism and absolutism, a precedent that will inevitably be discussed is Hetherington's gradualist proposal. As the most influential advocate of epistemic gradualism, Hetherington ushered many ways to motivate gradualism. Before I proceed any further to put forward my own arguments for gradualism, it will be helpful to analyse Hetherington's endeavours critically in this chapter. In a nutshell, Hetherington has developed two distinct versions of gradualist theories. Let us revisit them one by one.

1. The First Version

1.1: Introduction

In his 2001 book, Hetherington starts his defence of gradualism with a few linguistic evidence indicating that gradable expressions of 'know-that'

construction can be felicitous (see Hetherington 2001: 2-3). He argues that, although there are other (absolutist) ways to talk about knowledge, there is no good reason to deny the felicitousness of those gradable statements that he listed. Hetherington's explanation for why those statements are consistent and felicitous is: 'knowing a fact' is closely related to 'understanding a fact', and understanding is something clearly gradable. Hetherington argues that knowing a fact is sometimes equivalent to understanding that fact—understand that the fact obtains, and understand *how* that fact obtains. Hetherington also quotes Edward Craig's view that knowing-how is sometimes indistinguishable from knowing-that. Since knowing-how is gradable, so is knowing-that.

In addition, Hetherington also draws an analogy between knowing a fact and knowing a person. He argues that both sorts of knowledge can be gradable in the same manner: If one could only answer a few basic questions about the person (What is his name? What is his gender? How does he look like? *etc.*), one knows that person poorly; if one could answer every possible question about that person (What is his religious faith? What is his food preference? What is his favourite basketball team? *etc.*), one knows that person perfectly. Similarly, if one could only answer a few basic questions about a fact (Is it true? When did it happen? *etc.*), one knows that fact poorly; if one could do answer every possible question that might be raised about that fact (How to prove it? How does it obtain? Why does it obtain? *etc.*), one knows that fact perfectly. For Hetherington, knowing a fact and knowing a friend is fundamentally homologous, thus he notes that 'knowing is tantamount to becoming friends with a fact' (Hetherington 2001: 9).

Now we have seen a rough sketch of the first approach that Hetherington takes to argue for gradualism. That is, to draw analogies between knowing-that and relevant gradable concepts like understanding, knowing-how and knowing a person. Call it *the argument from analogy*. Apparently, this is a weak form of argument for motivating epistemic gradualism. It can help us to

see that the idea of gradualism might be more plausible than it is standardly assumed to be, but it can hardly prove that gradualism is correct in a rigorous way. I am not saying that this sort of argument from analogy is invalid, rather, I am just saying that this argument is weak. After all, talking of anything *analogous* to knowledge-that is not talking of knowledge-that *per se*. Moreover, Hetherington does not articulate *how* analogous knowledge-that is to those gradable epistemic notions. Except for the relation between how-knowledge and knowledge-that, which I will put weight on later, Hetherington does not seem to plan to advocate that knowledge-that is *tantamount to* those analogues. Thus it is also unclear how close the connection between the gradability of knowledge-that and that of its analogues is.

The second approach that Hetherington takes to argue for gradualism is to show *how knowledge can come in different degrees*. Call this *the argument from showing-how*. In his 2001 book, Hetherington proposed many theoretical frameworks within which knowledge-that can be understood as gradable. For example, Hetherington suggests that one's knowledge of a given proposition can be improved by virtue of *education*. A child's knowledge that 'there are kangaroos' can be improved after she became a biology PhD investigating kangaroos, because her knowledge of the existence of kangaroos has become deeper and more entrenched thanks to her prolonged education on this topic. Hetherington notes that it would be counterintuitive to insist that years of biological education can merely help her to maintain her childhood knowledge that there are kangaroos. If advanced education can improve one's knowledge of a given proposition, then one's knowledge-that-*p* can be graded along a qualitative scale. Hetherington also argues that one's deeper and more extensive *knowledge of a topic* can improve one's knowledge of a fact related to that topic. A biology PhD student's knowledge of the topic of kangaroos (kangaroos' habits and habitats, kangaroos' reproductive forms, *etc.*) can improve her knowledge that 'there are kangaroos'. 'Increased holism can improve your knowledge, even on a specific and simple fact' (Hetherington 2001:23). In

addition, it is also proposed that one's knowledge-that-*p* can be graded in terms of its failability (*nota bene*, not fallibility⁴⁰). A subject S knows that *p* failably, *iff*, (1) S knows that *p*; (2) there are possible worlds where (i) *p* is false; or (ii) S fails to believe that *p*; (iii) S fails to have good justification for *p*. Failability can be graded in terms of the amount and the distance of those 'failure-worlds' (*i.e.*, worlds characterised in (i), (ii), and (iii)). The more or the closer those failure-worlds are, the more one's knowledge that *p* is failable. Accordingly, the more failable one's knowledge that *p* is, the worse one knows that *p*. Given that the strength of justification usually determines failability, the corollary is: in general, the weaker one's justification for *p* is, the worse one's knowledge that *p* is.

Hetherington also applies this view to solve the Gettier problems. He names those epistemically lucky/risky factors involved in Gettier cases as 'strange occurrence'. Hetherington argues that despite those strange occurrences, subjects in Gettier cases still possess knowledge—but just know poorly. Hetherington remarks that:

'One's having a well-justified true belief that *p* is sufficient for one's knowing that *p*. But if there is a Strange Occurrence within one's context, the knowing is failable.' (Hetherington 2001:76)

And furthermore,

'We need to keep in mind that—other things being equal—the more failable a piece of knowing is, the less

⁴⁰ In brief, for Hetherington, 'knowing failably' is a broader concept than 'knowing fallibly' – the latter means that S knows that *p* but S's belief that *p* could have been *false*, while the former means that S knows that *p* but S could have *failed* to do so. Failability is broader than fallibility in that the former includes not only possible worlds where S's belief could be *false*, but also worlds where S fails to believe that *p*, and worlds where S fails to have good justification for *p*. Hetherington suggests that, compared with fallibility, failability is a preferable notion in the sense that it can take better care of knowledge of necessary truths – this sort of knowledge can fall short of being certain, while cannot possibly be fallible; in contrast, we can describe them as being failable. For a full discussion about the distinction between failability and fallibility, see Hetherington (1999).

confident we should be that it is knowledge.’
(Hetherington 2001:86)

A piece of completely infailable knowledge can be deemed as a piece of perfect knowledge. A very failable item of knowledge is a very poor item of knowledge. That is how knowledge can be graded in terms of failability. A natural question people will ask is: how poor a piece of knowledge can be qua knowledge? How failable can a piece of knowledge be to distinguish itself from non-knowledge? What can an item of *minimal knowledge* be like? Hetherington’s answer to these questions is, somewhat surprising and radical, that mere true belief can be sufficient to constitute the minimal knowledge. Justification is unnecessary for (the minimal) knowledge. That constitutes his objection to the so-called second ‘dogma’ (along with *epistemic absolutism*) of epistemology in his 2001 book, namely *justificationism*—the standard epistemological view that propositional knowledge entails justification. As radical and revolutionary this anti-justificationist view is, Hetherington (2001) employs a fairly weak argument to defend it. His argument for anti-justificationism that I will introduce next, as Hetherington admits, is a sort of argument from analogy as well.

Hetherington rejects justificationism by criticising both externalist (e.g., reliabilism) account and internalist (e.g., evidentialism) account of justification. He accuses them of both implausibly requiring that one has to know a topic if one is to know a fact. Hetherington argues that this reverses the real conceptual relationship between knowledge of a topic and knowledge of a fact—‘Knowledge of a topic is conceptually dependent on knowledge of facts—not vice versa’ (Hetherington 2001: 113). To hammer home this point, Hetherington invites readers to consider an exam case where an examinee Mario is asked to answer 100 multiple-choice questions. Hetherington claims that reliabilism requires one’s well-justified true belief to be repeatable. The idea is that in the exam case, for Mario’s knowing one question’s answer, reliabilism will expect Mario to be able to give many correct answers rather than only one correct answer. Alternatively, Mario

should be able to repeat the same true answer in similar possible situations. Otherwise, Hetherington claims that, reliabilism would not recognise Mario as knowing the correct answer. Hetherington thinks this repeatability requirement is too harsh. Here, the analogy he draws is that reliabilism sees knowing-that as something that is more akin to a *good act* done by a *good person*, rather than a good act *simpliciter*. Hetherington holds that the correct way to interpret knowing-that is to treat it just as a good act *simpliciter*. Just like a good act *per se* can be done by a bad person and can be one-off, a piece of knowledge-that-*p* can just be attained by virtue of some *unique characteristics* (some lucky factors involved) of one's process of attaining that piece of knowledge. For Hetherington, we cannot simply rule out the possibility that knowledge can also be one-off and thus unrepeatable. That is the analogy Hetherington trades on to argue against externalism. When it comes to internalism, his analogy becomes one between a supporting essay and supportive evidence. Hetherington claims that, for evidentialism, Maria's knowing a particular fact (knowing the answer to one question) depends on her having knowledge of the background topic. It is just like treating her supportive evidence to her knowledge of a fact as a supporting essay on the relevant topic. For Hetherington, evidentialism is ridiculously demanding as evidentialism is simply requiring that Maria has to be able to write a supporting essay on the relevant topic to show that she knows a particular fact within that topic.

One might find Hetherington's argument from analogy unconvincing. After all, it is far from clear how much analogousness there is between externalism/internalism and Hetherington's analogy. Also it is highly doubtful whether externalists and internalists would actually impose those (allegedly implausible) requirements to knowledge as Hetherington predicts they would do. Given the ambiguousness and doubtfulness of those analogies, I do not plan to give an intensive criticism to Hetherington's original argument for anti-justificationism here. However, it is noteworthy that, apart from his original arguments, Hetherington also refers to Alvin

Goldman's anti-justificationist view. Here is the argument that Hetherington quotes from Goldman:

'In one sense of "X knows that P ", it is synonymous with "X is aware that P " (or "X is apprised of P "), a sense that ignores justification. Suppose it is given that P is true, and we wonder whether Jane is aware of it. The only question that needs to be resolved is whether she believes P . If she does, she is aware of it; if she doesn't, she is unaware of it. The issue of justification or evidence is irrelevant. "Know" can be used similarly. If we wonder whether Jane knows that p , again given its truth, the only issue to be settled is whether she believes it. She knows if she does believe it, and is ignorant (does not know) if she does not believe it. The issue of justification, or its ilk, is again out of the picture. Here is another example. The sentence "You don't want to know what happened while you were gone" seems to mean: You don't want to have the truth about what happened in your belief corpus. It does not seem to require the translation: You don't want to have a *justified* belief in the truth of what happened. So I believe there is an ordinary sense of "know" in which it means "truly believe.'" (Goldman 1993: 24-25)

Hetherington accepts Goldman's verdicts that we sometimes use 'knows' without requiring justification. Nonetheless, it is hard to say that Goldman's argument for anti-justificationism is as convincing as Hetherington takes it to be. The quotation above basically consists of two arguments. The first one essentially turns on the linguistic resemblance between 'knows that p ' and 'be aware that p '. On the basis of this *prima facie* resemblance, it is concluded that one's knowing that p is equivalent to one's being aware that p . However, this equivalence will always be dubious before we have a clear definition of (epistemic) 'awareness' here, not to mention that it is also unclear why one's being aware that p does not include justification. A belief that suddenly pops up in one's mind out of nowhere without any support can hardly constitute knowledge, even if the belief is true. To attain knowledge that p , it is ordinarily expected that there is a certain sort of *connection* between one's belief of p and the truth of p . If knowing that p is equivalent to

being aware that p , and being aware that p just requires believing that p (as Goldman predicts), then where is that kind of *connection* between belief and truth?

Goldman's second argument relies on a seeming fact that the linguistic interpretation of some sentences involving 'knows' does not require justification. However, as we have seen in the last chapter, our linguistic uses of 'knows' can hardly accurately capture the conceptual nature of knowledge-that. That is because, we do not always use 'knows that' strictly in daily linguistic practice. Linguistic interpretations of those 'loose' uses of 'knows that' are inclined to be too flexible to reflect the genuine characteristics of knowledge-that. Some counterintuitive conclusions might also be derived by conducting linguistic analyses as Goldman does. Suppose that your friend Elvin spoke to you that 'you don't want to know what happened while you were gone'. And then he told you that after you left Maria's birthday party last night, Maria threw the birthday gift you gave her into the bin. Unbeknownst to you, Elvin just played a trick on you. Maria did not throw your gift into the bin. Elvin lied to you just in order to make you embarrassed. In that case, the proper translation of Elvin's statement 'you don't want to know what happened while you were gone' should be: you don't want to *hear* what I am going to tell you. It does not seem to require the translation: you don't want to have a *true belief* regarding what happened in your belief corpus. Whether the belief is true or not does not really matter, as what is unwelcome for you is not the *truth* about what happened after you left, but rather what Elvin *made up* about what happened after you left. Moreover, maybe Elvin does not even expect you to *believe* what he said—your embarrassed doubtful face is enough for Elvin's purpose. Hence, neither truth nor belief is required for interpreting Elvin's statement linguistically. Given these, it seems that Goldman's linguistically sensitive argument would also imply that knowledge does not entail true belief.

However, Hetherington's second approach to argue for gradualism, namely, the argument from showing-how, is still weak. After all, what Hetherington does is just constructing a competing theory to the standard absolutist theory. Hetherington frequently uses phrases like 'the standard view is not mandatory', 'it is not mandatory to think in the standard way', *etc.* However, what essential for gradualists is to prove the competing theory to be *preferable*, rather than just *conceivable*. In other words, to prove the standard view to be *not satisfactory*, rather than just *not mandatory*. Merely showing how a theory can be constructed does not suffice to achieve this. After all, an extreme case is, you can even attempt to prove that the world is made of hamburger by showing how it is possible that this world is made of hamburgers if we endorse an alternative definition of 'hamburger'. But arguments like this is far from convincing. Surely, I am not saying Hetherington's argument from showing-how is as ridiculous as 'hamburgerism'. In fact, in some sense, the way in which Hetherington constructs his gradualism is fairly promising. It is just not a strong argument for epistemic gradualism. It is more like a description of what knowledge would be like, *if* epistemic gradualism is true. But the central issue on the table is why epistemic gradualism is true.

Before ending this subsection, the last point that I want to emphasise is that there is a seemingly close connection between Hetherington's *argument from showing-how* and his anti-justificationist view. Hetherington takes anti-justificationism as a benefit that his gradualism can bring about. That is, justificationists will be troubled by the so-called 'boundary problem' of justification. The boundary problem asks: where is the cut-off point among the spectrum of justification that distinguishes *sufficient* justification constituting knowledge and *insufficient* justification that fails to constitute knowledge? Hetherington thinks that it is hard to find such a non-arbitrary cut-off point, while discarding justificationism can avoid this issue—mere true belief can suffice to constitute the minimal knowledge. In light of this, Hetherington takes 'gradualism + anti-justificationism' to be superior to

'absolutism + justificationism'. Thus, Hetherington might construct a strong argument for gradualism if he can successfully prove two points: 1) anti-justificationism does enjoy overall advantages over justificationism, and 2) anti-justificationism is intrinsic to gradualism and justificationism is a mandatory choice for absolutism. Notwithstanding, in what follows, I will prove that both points are untenable. I do not think epistemic gradualism will perforce lead one to embrace anti-justificationism as Hetherington does. This will be clearer as my argument unfolds. Also, I do not think anti-justificationism is preferable to justificationism—I will hammer home this point in section 3.

1.2: Discussion

Three salient characteristics of Hetherington's first version of epistemic gradualism can be summarised as follows:

- (1) It *leads to* (if not being based on) anti-justificationism.
- (2) It is a combination of threshold *absolutism* and beyond-threshold gradualism.
- (3) It implies that there should be an absolute definition of the minimal knowledge.

The first characteristic has been introduced above. In endorsing an anti-justificationist view, Hetherington proposes an account of knowledge as follows:

'For any epistemic subject x , time t , and proposition p :
(1) At t x knows that p =*df.* At t , x has a justified and true belief that p ; (2) this knowledge that p is better, or it is worse, as knowledge that p to the extent (respectively) that x 's justification for this true belief that p is better or worse.' (Hetherington 2001: 90)

Justification is not necessary for this account of knowledge. However, Hetherington also concedes that justification is required if one wants to know *better* that *p*, rather than to just attain the minimal knowledge. More justification can improve one's knowledge. Moreover, every single item of justification can improve one's knowledge—as opposed to the beyond-threshold absolutist view that propositional knowledge cannot be improved by more or better justification once the threshold of knowledge is met. This brings us to the second trait of Hetherington's first version of gradualism. Hetherington claims that:

'In denying that knowledge need be absolute (as I will be doing in this chapter and beyond), I will not be denying that there is a cut-off point—and an absolute one at that—between knowing and not knowing. When 'Is it really knowledge?' is used like 'Is it really a toy?', it does not commit the speaker to there not being better or worse knowledge of a particular fact. Maybe she is unsure as to whether some specific belief falls into the category of being knowledge. But even that uncertainty is compatible with her thinking that there can be both better and worse cases of knowledge qua knowledge—more carefully, cases of knowledge that are more, and ones that are less, clearly cases of knowledge, more or less clearly members of the category of being knowledge.' (Hetherington 2001:6)

That is to say, Hetherington claims that he is committed to a combination of views as follows:

'(i) Knowledge is to be absolutely distinguished from whatever is not knowledge (whatever is external to knowledge). This is because there is an absolute cut-off point between knowing and not knowing, (ii) But *within* the category of knowing, non-absolutism is true. That is because it is possible that some cases of knowledge that *p* are better as knowledge that *p*, or more clearly knowledge that *p*, than other cases of knowledge that *p*.' (Hetherington 2001:7)

View (i) is called 'external absolutism' by Hetherington (and in my terminology, 'threshold absolutism'), and view (ii) is named 'internal non-absolutism' (in my terminology, 'beyond-threshold gradualism'). I do not think Hetherington can successfully develop epistemic gradualism by defending this combination of views. That is because, when endorsing the external absolutism and pursuing an internal non-absolutism, Hetherington confounds the *fuzziness* of knowledge with the *gradability* of knowledge. He can at most prove that 'knows' is fuzzy, rather than gradable. The claim that 'X is ungradable' can be compatible with the claim that 'there are less or more clearly cases of X *qua* X'. For instance, 'death' is not a gradable concept, it will be untoward to say 'A is more dead than B is'. Nevertheless, there can also be less clearly cases/manifestations of death (e.g., lying motionlessly without breath) and more clearly cases/manifestations of death (e.g., brain death indicated by the electroencephalogram). The fuzziness of the boundary of a concept X does not suffice to entail the gradability of X. After all, it can be argued that every philosophically interesting concept is fuzzy in this or that sense.

To be more specific, there can be at least two kinds of fuzziness: first, the *definitional fuzziness*, which means that the boundary of a concept X is fuzzy due to the vagueness of the definition of that concept—due to the lack of an absolute criterion to determine whether something can be recognised as an instance of X or not. For example, the fuzziness of being a toy. It is hard to define what a toy is in an absolute and clear manner. Sometimes we adopt functional definitions, e.g., 'a toy is an artefact that can be played with', but there is no absolute criterion to judge whether an object can be played with or not. Furthermore, it is even hard to define the behaviour 'play with' precisely. The second kind of fuzziness, which I shall call the *epistemic fuzziness*, means that a concept X has a clear and absolute definition, but sometimes it is not easy to tell its instances in our daily epistemic practice. For example, we can have an absolute definition of 'fake' (counterfeited; not real), but sometimes it is still hard for us to tell whether a given artwork is fake or not.

Also, sometimes it is easy to tell that a poorly counterfeited fake is fake. Hence, there are still less or more borderline cases of 'fake' even though the definition of 'fake' can be absolute. Notwithstanding this fuzziness, 'fake' is still an ungradable term. The fuzziness of 'death', I think, also falls into the category of epistemic fuzziness.

The definitional fuzziness often leads to the epistemic fuzziness, but *not* vice versa. Among the two sorts of fuzziness, the epistemic fuzziness of concepts is much more clearly compatible with the ungradability of concepts. It seems to be relatively hard to find a predicate that has a vague definition but is still ungradable. But it is easy to find predicates that are epistemically fuzzy but still ungradable (e.g., 'fake', 'dead', 'extinct', 'plagiaristic'). One explanation for this difference might be: almost every concept, under certain conditions, can be epistemically fuzzy; while not every concept is definitionally fuzzy (admittedly, the concept 'fuzzy' itself is definitionally fuzzy and gradable, but it does not mean that everything has a fuzzy definition). We will revisit the relation between vagueness and gradability in the next chapter when we remodel the debate between gradualism and absolutism. At this stage, let us move forward to see how this relates to Hetherington's gradualism.

Given that Hetherington endorses that there is an absolute difference between knowledge and non-knowledge, it seemingly implies that Hetherington agrees that the definition of knowledge is clear and absolute so that there can be a cut-off point of knowledge. As a result, it seems that for Hetherington, the fuzziness of 'knows' is not definitional but only epistemic. The reason why there are less or more clearly cases of knowledge *qua* knowledge is not because of the vagueness of the definition of knowledge, but because of the practical difficulty of identifying instances of knowledge, the practical difficulty of making knowledge ascriptions. However, we have noted that epistemic fuzziness is insufficient to entail gradability. Therefore the basis of Hetherington's internal non-absolutism is

not solid. In addition, since every concept can be epistemically fuzzy under certain conditions, it is hard to see who the opponent of Hetherington's internal non-absolutism will be, if this view is just based on the sense that the concept 'knows' is of epistemic fuzziness. That is, internal non-absolutism may not be that controversial as Hetherington presents it to be, and on the contrary, internal absolutism is not that platitudinous as Hetherington takes to be.

Now we have discussed characteristics (1) and (2) of Hetherington's first gradualism. With (1) and (2) in play, a natural corollary that would be entailed is (3): that there needs to be an absolute definition of the minimal knowledge. That is because, (1) implies that justification is not necessary for knowledge, and (2) claims that there is an *absolute difference* between knowledge and non-knowledge. Combining (1) and (2) together, the absolute difference between knowledge and non-knowledge just amounts to the absolute definition of the minimal knowledge, *i.e.*, true belief. The definition of the minimal knowledge has to be absolute (without definitional fuzziness) so that the difference between knowledge and non-knowledge can be absolute—a straightforward definition fitting the bill is 'true belief'.

I will reject Hetherington's first version of gradualism in terms of the three characteristics. Precisely put, I do not accept anti-justificationism; I do not accept the threshold/external absolutism that Hetherington accepts (I deny that there is an absolute difference between knowledge and non-knowledge); I do not think an absolute definition of the minimal knowledge is necessary.

There is more that needs to be said about the relation between Hetherington's minimalism and his gradualist proposal. Strictly speaking, Hetherington's gradualism is not logically *based on* his minimalism—it is not the case that the former entails the latter or the latter entails the former. However, the two views are closely connected to each other. As we have

noted, Hetherington's first version of gradualist proposal is a hybrid one consisting of internal gradualism and external absolutism—he insists that there is a cut-off boundary distinguishing knowledge from not-knowledge. Correspondingly, now that knowledge is argued to be gradable and there can be good knowledge and bad knowledge, a natural question to be asked is how *bad* an instance of knowledge is allowed to be. In other words, where is the *boundary* of knowledge⁴¹? Given his internal gradualist and external absolutist stances, minimalism turns out to be a straightforward answer for Hetherington. Again, this minimalist answer is not the only logically possible option (and might also not be the optimal option) for Hetherington, but it does constitute an unignorable part of the whole gradualist storytelling as a matter of fact. It is motivated by Hetherington's commitment to external absolutism, and can, in turn, answer an important (at least for Hetherington) problem invited by this commitment. Therefore, minimalism completes Hetherington's first version of gradualist picture as an important piece of jigsaw. I concede that rejecting Hetherington's minimalism is insufficient for rejecting his gradualism. Nevertheless, the defects of his minimalism should also be seen as defects of his gradualist account of knowledge.

I have outlined the first version of Hetherington's gradualist proposal, now let us go ahead to see how Hetherington constructs his second version of gradualism.

2. The Second Version

2.1: Introduction

⁴¹ This seems to be a very important philosophical issue for Hetherington, and he has discussed this problem in many places (e.g., 2001. Chap 4; 2006; 2011. Chap 1). Notice that, this is not an exclusive problem for gradualism. Hetherington (2006) argues that anyone who endorses fallibilism should pay adequate attention to the boundary problem. However, as we will see in Chapter 4, if we reject external absolutism, then gradualists do not have to be nagged by the boundary problem too much.

Hetherington's second version of gradualism originates from his 2005 paper 'Knowing (How It Is) That P: Degrees and Qualities of Knowledge', where knowing that p is argued to be identical with *how-knowledge that p* , to wit, knowing how it is that p . This idea is inherited and expanded in his 2011 book 'How to Know' (in particular, the fifth chapter).

Before we proceed to introduce the main argument for Hetherington's second version of gradualism (which is based on how knowledge-that p), it is worth noticing that Hetherington (2011) also puts forward a closely related but subtly distinct approach to argue for gradualism, which is motivated by a novel account of knowledge that he calls *practicalism*. Hetherington's practicalism, as the reverse of the standard intellectualism, argues that knowledge-that is a kind of knowledge-how. One core practicalist thesis relevant here is the so-called *knowledge-as-ability hypothesis*, which interprets knowledge as an ability to manifest various accurate representations of p . This ability can be graded in accordance with how detailed one knows how it is that p . The more aspects of how it is that p one knows, the better one's knowledge that p is. Accordingly, knowledge, *qua* an ability, is gradational.

There are three reasons why I do not see Hetherington's practicalist account of gradualism as the core of his second version of gradualism. First, Hetherington's practicalist argument is not as exclusively dedicated to gradualism as his argument based on how-knowledge that p that we will introduce later. The main target the knowledge-as-ability hypothesis aims to argue against is the standard view that knowledge is a kind of belief, rather than that knowledge does not admit of degrees. Gradualism only serves as a 'side product' if the knowledge-as-ability hypothesis succeeds⁴². Second, I

⁴² I have qualms about the hypothesis itself as well. The hypothesis argues that knowledge that p is the ability to register accurately that p . Therefore, knowledge is not the manifestation of abilities, but the ability itself. On the contrary, true belief is just a type of manifestation of knowledge, while knowledge *per se* is the ability to register accurately that p . 'Accurately believing that p ' is just a member of what Hetherington calls *epistemic diaspora*, which refers to a group of manifestations of knowledge including believing

think the account of gradualism based on how-knowledge that p that Hetherington presented in the fifth chapter of the same book is a more robust expansion of the practicalist account that he presented in the second chapter. Although the two accounts appeal to different conceptual resources to construct gradualism, these resources are closely connected to each other. For example, both accounts rest on an essential notion ‘how it is that p ’. The practicalist account of gradualism suggests that knowledge that p , *qua* an ability, can be graded in accordance with how detailed one knows how it is that p . How should we understand ‘how detailed one knows how it is that p ’? I think, the account of gradualism based on how-knowledge that p can give us a more explicit answer. That is, ‘how many details of how it is that p are known’ can be interpreted in terms of ‘how many parts of a minimal truthmaker for p are known’. Finally, even if the knowledge-as-ability hypothesis is correct, it is unclear how far this can lead us towards gradualism. Take Hetherington’s own ‘home run’ analogy as an example. Hetherington (2011) concedes that striking a home run is an ability, however, whether a home run is struck or not is an absolute yes-or-no affair. He also notes that ‘knowledge—according to knowledge-absolutism—is like that (home run)’ (2011: 7)⁴³. If that is how absolutism interprets knowledge (in fact, this is a typical external-absolutist interpretation), then it will be confusing how knowledge’s being an ability helps to refute absolutism. Striking a home run is an ability which can be evaluated and graded by many qualitative properties, such as the amount by which the ball clears the outfield wall or the speed at which it flies. But whether a home run is struck or

accurately that p , asserting accurately that p , reasoning accurately in p -related ways, acting appropriately in p -related ways, *etc.* However, it is unclear why the existence of the *epistemic diaspora* entails that knowledge is an ability. Why is it not the contrary? Why is it not the case that those members of the epistemic diaspora constitute what we call ‘cognitive competence/ability’ and knowledge is just the *manifestation* of this competence/ability rather than the competence/ability itself?

⁴³ This reading of epistemic absolutism is borne out by Fantl & McGrath (2009: 24): ‘Home runs depend on, and vary with respect to, factors that come in degrees: how far the ball traveled, how hard it was hit, how skillfully, *etc.* There are, in some sense, better and worse home runs—home runs that are more or less majestic, more or less lucky, *etc.* But whether a hit is a home run is a binary matter. There is no reason the fallibilist shouldn’t say the same thing about knowledge.’

not is an absolute binary matter. Similarly, shooting a three-pointer is an ability, but 'shooting a three-pointer' is not gradable. Thus it seems that ' ϕ is an ability' fails to entail that ' ϕ is gradable'. At least, this type of ability-based gradualism would still allow for the equivocal attitude that I mentioned in Chapter 1: knowing that p , *qua* an ability, can be graded in a peculiar sense—but it is ultimately non-gradational in general, just like hitting a home run and shooting a three-pointer.

Comparatively, Hetherington's how-knowledge-that- p -based account of gradualism is more fully developed. At the start of his 2005 paper on this form of gradualism, Hetherington states in the footnote that:

'Unlike the knowledge-gradualism in that earlier book, this new version will not be framed in terms of whatever strength of justification is present within a given piece of knowledge.' (Hetherington 2005: 129-130)

Indeed, in the beyond-threshold gradualism that Hetherington develops previously, what determines the degree of a piece of knowledge is the strength of its justificatory support. The better justification one has, the better her knowledge (beyond the threshold of knowledge) will be. While in his new account of gradualism, Hetherington resorts to the idea of how-knowledge that p ⁴⁴ to explicate how knowledge-that p can be gradable:

'That is, in my previous knowledge-gradualism, the degree or quality of a piece of knowledge was a function of what constituted the knowledge 'from within'—the strength of its justificatory component. In this paper's knowledge-gradualism, the degree or quality of a piece of knowledge that p is a function of the extensiveness of the 'wider' network of knowledge that constitutes the degree to which the knowledge that p how-knows that p .' (Hetherington 2005: 138)

⁴⁴ How-knowledge that p , Hetherington emphasises, is distinct from knowledge-how. The former refers to knowledge of how it is that p , whereas the latter refers to knowledge of how to do something.

Hetherington does not provide a very clear definition for how-knowledge. Basically, he uses this term to refer to *knowledge of how it is that p*. And *how it is that p*, according to Hetherington, comes in many aspects. The minimal aspect of *how it is that p* is that *p obtains*. That is to say, the minimal aspect of one's how-knowledge that *p* is one's knowledge that *p simpliciter*. How-knowledge that *p* can be enriched and deepened if more aspects of how it is that *p* (other than that '*p obtains*' *per se*) is known.

Hetherington further trades on the truthmaker theory to spell out the notion of how-knowledge. Hetherington adopts David Armstrong's truthmaker necessitarianism (see Armstrong 2004) and defines a truthmaker for a proposition *p* as something in the world that necessitates *p*'s truth being true. A *minimal* truthmaker for a truth that *p* is characterised as a truthmaker such that if some parts of it are subtracted, then it will not suffice to make it true that *p*. For example, a minimal truthmaker for the truth that 'Kyle is a boy' can be Kyle's being a boy. Given these, to have a better how-knowledge that *p* is to know more parts of a minimal truth-maker for *p*. Hetherington thereby proposes a thesis 'HkTm':

'HkTm At time *t*, a person *x*'s how-knowledge that *p* is more, or it is less, well developed or extensive, in accord with how many parts of a minimal truthmaker for *p* are known by *x* at *t*.' (2005: 137)

So how does Hetherington utilise those concepts to argue for his new epistemic gradualism? The schema of his argument can be summarised as follows, which I refer to as *argument from equivalence*:

[Argument from Equivalence]

(E1) Having knowledge that *p* entails having how-knowledge that *p*, as knowledge that *p* is the minimal aspect of how-knowledge that *p*.

- (E2) Having how-knowledge that p entails having knowledge that p , as knowledge that p is the minimal aspect of how-knowledge that p .
- (E3) From E1 and E2, having knowledge that p is the necessary and sufficient condition of having how-knowledge that p .
- (E4) From E3, having knowledge that p is equivalent to having how-knowledge that p .
- (E5) How-knowledge that p is gradable.
- (E6) Knowledge that p is gradable.

Premises E1 and E2 are taken to be true by virtue of Hetherington's characterisations of how-knowledge that p . It is stipulated that the minimal requirement of one's knowing how it that p is knowing that p *simpliciter*. Knowledge that p is deemed to be the threshold of how-knowledge that p . Therefore knowing that p entails having how-knowledge that p (to the minimal extent), and vice versa. From E1 and E2, it seems that one can derive E3 and E4 smoothly. Hetherington presents E3 as a biconditional that he calls ' $\leftrightarrow H$ ', to wit, knowing that p both entails and is entailed from knowing how it is that p . From that, Hetherington infers that 'knowledge that p is how-knowledge that p ' (Hetherington 2011: 177). Premise E5 also seems to be apparently true by lights of HkTm—one's how-knowledge that p can be graded in terms of how many aspects of how it is that p is known.

One's knowing how Kyle is a boy might be minimal if only that fact of Kyle's being a boy is known. It can be improved if one also knows the age of Kyle, or knows that Kyle is a student of a boy's school, *etc*. If E5 is true, and knowledge that p is equivalent to how-knowledge that p , then knowledge that p should also be gradable—E6 can thus be inferred.

2.2: Discussion

Before we proceed to see how Hetherington's argument from equivalence is untenable (which will be done in section 4), let me end this chapter by summarising some traits of Hetherington's second version of gradualism. In particular, we shall focus on places where it differs from the previous version.

First of all, it is obvious that Hetherington's second version of gradualism is based on the concept of how-knowledge that p and the truthmaker theory, which is quite different from his first version of gradualism which is based on the analysis of epistemic justification. Consequently, the distinction between beyond-threshold gradualism and threshold gradualism becomes marginal in his second version gradualism, as this distinction no longer plays as important a role as it does for the first version of gradualism.

As a result, Hetherington also characterises his new concept of 'minimal knowledge' in a different way from his previous version. Various interpretations of his new account of 'minimal knowledge' were given in his 2005 paper:

'When a person x 's how-knowledge that p is restricted to knowing a minimal truthmaker for p , x has *minimal* knowledge that p ' (Hetherington 2005: 138)

And,

'[M]inimal knowledge that p is unaccompanied by knowledge of any parts of a minimal truthmaker for p —any parts beyond p itself.' (Hetherington 2005: 141-142)

Moreover, Hetherington also identifies *minimal knowledge* with *fundamental knowledge*:

'A piece of foundational knowledge needs only to be an instance of knowledge that p which is not informative (and draws upon no information) about anything other than p .' (Hetherington 2005: 143)

Here, Hetherington's definition of the minimal knowledge differs significantly from his previous work. Minimal knowledge is not interpreted in terms of justification, but rather in terms of minimal truthmaker. An item of minimal knowledge, here, is no longer understood as one that supported by no justification, but rather one that accompanied by no other aspect of how-it-is-that-*p* other than *p per se*. Just as Hetherington emphasises, his previous understanding of minimal knowledge is 'from within'—focusing on the justificatory components within a given piece of knowledge; whereas in his new version of gradualism, a holistic understanding of knowledge is employed—now he focuses on the holistic connection between a given piece of knowledge-that-*p* and its related supportive how-knowledge-that-*p*⁴⁵. Hetherington admits this change and claims that:

*'How unimpressive is minimal knowledge? For example, is it mere true belief—a true belief bereft of any justificatory support? I am not presuming so. Hence, I am explicating its minimality or bareness in terms of its object (namely, *p*—and nothing else), rather than in terms of an absence of justificatory support for it. I will continue leaving open that question of whether cases of minimal knowledge enjoy justificatory support.'*
(Hetherington 2005: 142)

Again, we can spot another important difference between Hetherington's two versions of gradualism: the first version implies that knowledge does not entail justification, that is, the first version implies anti-justificationism; while the second version weakens its advocative attitude towards anti-justificationism—Hetherington leaves the question of whether justification is necessary for knowledge as open⁴⁶. Hetherington's second version of gradualism can thus be deemed to be less radical than the previous version. Notwithstanding, in what follows, it will be argued that both versions of Hetherington's gradualism are problematic.

⁴⁵ In fact, this holism concept was already mentioned in his 2001 book, where he contributed this idea to W.V.O Quine. However, at that time this holist concept of knowledge did not play as essential a role as it does in his second version of gradualism.

⁴⁶ Hetherington (2011, Ch 4) provides a more detailed discussion on this issue.

3. Criticism of Hetherington's First Gradualist Theory

In section 1, we have noted that there are three traits of Hetherington's first version of epistemic gradualism:

- (1) It leads to anti-justificationism.
- (2) It is a combination of threshold *absolutism* and beyond-threshold gradualism.
- (3) It implies that there should be an absolute definition of the minimal knowledge.

I disagree with Hetherington in terms of the three traits. Trait 2 and trait 3 are closely related—that is because Hetherington maintains that there has to be a clear cut-off point between the minimal knowledge and non-knowledge. He advocates the threshold absolutism and the 'mere true belief' account of minimal knowledge. That is to say, his minimalism proposal is also rooted in his threshold absolutist inclination. In what follows, I will show why the three traits turn out to be three defects of Hetherington's first version of epistemic gradualism. Since trait 2 and 3 share the same threshold absolutism ground, this section will proceed by criticising Hetherington's anti-justificationist and threshold absolutist inclinations.

3.1: Anti-Justificationism

It is epistemologically orthodox that knowledge requires justification⁴⁷, especially when justification is understood as a sort of generic epistemic support (*e.g.*, epistemic warrant, as it is sometimes known). However, Hetherington argues against this orthodoxy and urges that

⁴⁷ For exceptions, see Sartwell (1991, 1992), Skidmore (1993), Goldman (1993), Hawthorne (2002), Kornblith (2008). For helpful criticisms to anti-justificationism, see Le Morvan (2002, 2005).

knowledge can be attained without justification. His main motivations for advocating anti-justificationism are based on two observations:

First, Hetherington finds that both mainstream externalist and internalist accounts of justification reverse the proper conceptual-priority-relation between 'knowing a fact' and 'knowing a topic'. He holds that knowing a fact should be conceptually prior to knowing a topic—that is, to know a fact, one does not need to have possessed knowledge of an associated topic already. On the contrary, knowing a topic is conceptually dependent on knowing facts, as to know a topic is to know many facts about that topic. Hetherington argues that both reliabilism (a prominent externalist theory) and evidentialism (a representative internalist theory) read the conceptual relationship in the wrong order—they take knowing a fact to be conceptually dependent on knowing a topic. Take reliabilism for example. Hetherington uses a case of exam knowledge to argue that, for a student's knowing the correct answer of a specific question, reliabilists would require the student to be able to give many *more* correct answers in sufficiently similar circumstances. That is, to know a fact, Maria is required to have knowledge of an associated topic. Hetherington holds that this reliabilist requirement is unacceptable. Instead, 'enough further knowledge to amount to her having knowledge of a related topic is not conceptually necessary for her knowing a particular fact within that topic' (Hetherington 2001: 113).

Second, Hetherington worries that if we take justification as a necessary condition for knowledge, then it will be difficult to solve the so-called *boundary problem*. That is, if justification is necessary for knowledge, then how much justification can suffice to yield knowledge? Where is the boundary between 'enough justification' and 'insufficient justification'? Hetherington is sceptical about whether the cut-off point can be found in a non-arbitrary manner. On another hand, Hetherington holds that there has to be a cut-off point between knowledge and non-knowledge. Putting all these considerations together, he infers that justification should not be seen as necessary, so that the boundary

problem can be solved, and the cut-off point between knowledge and non-knowledge can be maintained. The cut-off point, Hetherington argues, is 'true belief'. Contra most epistemologists, Hetherington argues that mere true belief can suffice to constitute *the minimal knowledge* (call this view *minimalism*).

Both observations cannot reasonably motivate anti-justificationism. In the previous chapter, we have seen that Hetherington's criticism of externalism and internalism can hardly lend strong support to anti-justificationism. That is because, firstly, there are far more accounts of justification other than reliabilism and evidentialism, but Hetherington's criticism only cast suspicion upon the two selected accounts and thus appears to be gerrymandered. Secondly, reliabilism and evidentialism can deny that they actually require that one has to already know a topic before she could know a fact. For example, when it comes to the case of exam knowledge, reliabilists can agree with Hetherington's assertion that:

'[N]o instance of her having knowledge, within the exam, of an individual fact depends conceptually on its being justified in a way that amounts to her having knowledge of an associated topic.' (Hetherington 2001: 113)

That is because reliabilists can explicate that they *do not* expect Maria's knowledge of a fact to be justified in a way that amounts to her *having* knowledge of an associated topic. Rather, they just require Maria's knowledge to be justified in a way that amounts to her *being able to have* knowledge of an associated topic. The reliabilist requirement can just be dispositional. Consequently, this can avoid Hetherington's criticism that one's knowledge of a fact has to depend conceptually on her knowledge of an associated topic. Dependency implies a certain sort of priority (temporal, causal or constituent; *etc.*). That is, '*A depends on B*' implies that B enjoys a certain kind of priority over A. Certainly, reliabilism does not require that one must have already possessed knowledge of an associated topic (temporally) *before* one knowing a fact. Also, reliabilism could deny that one's knowing a

topic is causally prior to one's knowing a fact. It is not the case that one knows a fact *by virtue of* her knowledge of the relevant topic. Rather, reliabilists can argue that one's true belief is well justified *by virtue of* the reliable belief-forming method, which does not guarantee that one has already possessed knowledge of the relevant topic, but just that one would be in a very good position to know relevant topic if the method were taken. Furthermore, it is also not the case that one's knowing a topic should be seen as a more basic concept than one's knowing a fact—reliabilism does not claim that one's knowledge of a topic constitutes one's knowledge of a fact. Hence, the reliabilist account of justification does not necessarily imply that knowing a topic enjoy any sort of priority over knowing a fact. Therefore, Hetherington's censure that 'reliabilism renders knowing a fact *conceptually dependent* on knowing a topic' seems to be unfairly arbitrary.

Now let us move forward to the boundary problem. In that respect, Hetherington worries that it is difficult (almost amount to impossible) to find out the cut-off point between sufficient justification (for constituting knowledge) and insufficient knowledge. As a matter of fact, this worry does not suffice to entail that we should abandon justification as a necessary constituent of knowledge. On the contrary, insofar as we see propositional knowledge as a gradable concept, then the boundary problem should not discourage us, as most gradable concepts have difficulties in nominating the cut-off point. Belief is gradable, belief needs credence. It is difficult to tell how much credence is sufficient for constituting a belief. However, this does not mean that credence is unnecessary for belief. Humour is gradable, humour requires 'being funny'. It is difficult to tell how funny is funny enough for being humorous, but this does not imply that humour does not need being funny. The difficulty in finding the cut-off point is not a problem, but a mark of a gradable concept. This also applies to propositional knowledge.

We have seen that Hetherington's motivations for rejecting justificationism are ill-grounded. So, why cannot his anti-justificationism be accepted? What is

the problem of endorsing mere true belief as the minimal knowledge? A straightforward answer is that it violates the widely-shared intuition that knowledge is incompatible with epistemic luck, at least, intervening luck⁴⁸; and the intuition that knowledge should be a kind of cognitive achievement (see Pritchard 2007, 2012a; Sosa 2007; Greco 2003; Carter & Pritchard 2013). One can luckily form a true belief *via* mere guessing or wishful thinking. But a true belief formed in these ways can hardly be admitted as knowledge as it is true by virtue of mere epistemic luck. Intuitively, if one's true belief B is to be counted as an item of knowledge, then there should be some sort of non-lucky relation between the formation of B and B's being true. For example, one forms the belief B and B would not have been easily false (*cf.* the 'safety' condition of knowledge, *a la* Pritchard). Or, one's belief B is true because it manifests one's relevant cognitive competence (*cf.* the idea of 'aptness', *a la* Sosa). An instance of unjustified true belief lacks this sort of non-lucky relation, and thus fails to constitute knowledge. Numerous relevant discussions have been undertaken in this fashion⁴⁹, so I do not plan to expand it too much here. Apart from this, I would like to discuss another

⁴⁸ Dialectically speaking, it is noteworthy that Hetherington (2013) has explicitly argued against the anti-luck intuition. His objection to the anti-luck condition of knowledge is based on his objection to the safety theory of knowledge. Roughly speaking, Hetherington argues that the safety theory succumbs to a modal fallacy. He asserts that advocates of the safety theory cannot show a *true belief* B falls short of knowledge due to its unsafety. That is because, to show a *true belief* B to be unsafe, one has to re-identify B in a close possible world W where B (formed via the same method) is false, however, one can only identify a *false belief* B with the same content in W, rather than the same piece of *true belief* B. Thus one cannot show a *true belief* to be unsafe. This criticism of the anti-luck intuition strikes me as a very unconvincing one. Firstly, it only focuses on the modal reading of the safety-account of epistemic luck, but as we will see in Chapter Five, there are other (perhaps more plausible) ways to account for epistemic luck (*e.g.*, Goldberg's epistemic entitlement account of luck). Secondly, Hetherington seems to distort the anti-luck intuition and the safety condition of knowledge. It is unfair to accuse the safety theory of conducting a modal fallacy because it has to appeal to a false (and thus different) piece of belief across worlds. That is because, safety, by its nature, is a subjunctive/counterfactual concept – a piece of true belief is unsafe if in sufficiently similar possible worlds its counterparts (with the same content and is formed via the same method by the same subject) would be false. Consider a similar modal concept 'unnecessarily'. A true belief is unnecessarily true iff in some possible worlds its counterparts are false. It is unreasonable to reject the idea of 'unnecessarily' by complaining that to identify a *true belief* to be unnecessarily true, one can only re-identify a *false belief*, rather than the same *true belief*, in a possible world. It is also untenable to conclude that therefore we cannot show a true belief to be unnecessarily true.

⁴⁹ For a helpful review of the two intuitions (anti-luck and cognitive achievement), see Pritchard (2012a).

drawback to Hetherington's minimalism. That involves *the value problem of knowledge*.

It is ordinarily granted by epistemologists that knowledge is more valuable than mere true belief⁵⁰. This tradition can be traced back to Plato's *Meno*, where Socrates raises the so-called 'primary value problem', *i.e.*, why knowledge is more valuable than true belief. Accordingly, there is also a secondary value problem, which concerns the issue of why knowledge is more valuable than any proper subset of its parts. These questions invite heated discussions in the literature of epistemology (see Jones 1997; Kvanvig 2003, 2005; Pritchard 2007; Goldman and Olsson 2009; *etc.*), and the mainstream view holds that the value problem of knowledge is a genuine problem⁵¹ in the sense that knowledge is preferable to mere true belief indeed—there is discrepancy in value to explain.

Hetherington's minimalism might also grant that the value problem is a genuine problem worthy of considerations. Even if true belief is the minimal knowledge, there can still be differences in value between mere true belief and more robust knowledge—however, this discrepancy is just a matter of degree. Hetherington often describes mere true belief as to the 'worst possible knowledge'—the least valuable knowledge. Hence, in response to the value problem of knowledge, Hetherington's minimalism might argue that the difference in value between mere true belief and knowledge lies in the justificatory strength. A better-justified item of knowledge is a more valuable item of knowledge. While a mere true belief lacks justification, and thus will be the least valuable instance of knowledge:

'Thus, if you have justification for a true belief that *p*, this improves something about your having that true belief; and if we let the true belief be knowledge already ..., then we

⁵⁰ As an exception, Kvanvig (2003) famously argues that knowledge is not distinctively more valuable than mere true belief. For responses to Kvanvig's challenge, see Greco (2009, 2010), DePaul (2009), Ryan (2017), *etc.*

⁵¹ For some exceptions, see Sartwell (1991, 1992); Goldman (1993).

may interpret the justification as making the knowledge better knowledge—yet without the justification's having been necessary to the knowledge's very existence.' (Hetherington 2001: 132)

Therefore, in accordance with Hetherington's minimalism, anti-justificationism can be compatible with the primary and secondary value problem of knowledge. Knowledge does often (but not always, given that for Hetherington sometimes knowledge is just mere true belief) enjoy more value, but this advantage is just a matter of degree—in terms of how well-justified an item of knowledge is (e.g., how fallible the justification is).

Nevertheless, Hetherington's minimalism will be in trouble when it comes to *the tertiary value problem*, which holds that 'the difference in value between knowledge and that which falls short of knowledge must be one of kind and not merely degree' (Pritchard 2007:104) and then asks what *kind* of distinctive value knowledge has. This tertiary value problem is taken to be significant, as it captures the sense that we usually see knowledge as *distinctively* valuable. If the difference in value between knowledge and mere true belief is just a matter of degree, then it would be unclear why it is knowledge that is of distinctive epistemic value for us. How can knowledge be seen as the central kind of concept in epistemological research that is of unique interest to us? The tertiary value problem requires us to find the *difference in kind* such that we can explain why an *excellent* item of mere true belief cannot be seen as epistemologically valuable and desirable as knowledge, at least, a poor instance of knowledge. On the contrary, Hetherington's anti-justificationism (and thus his minimalism) asserts that the former is sometimes as valuable as the latter indeed. If anti-justificationism is right, then knowledge only has a greater *degree* rather than a different *kind* of value than mere true belief, since mere true belief is counted as a *kind* of knowledge. After all, the worst possible knowledge is still a *kind* of knowledge rather than anything that falls short of knowledge. That is to say, the difference in how well-justified a true belief is cannot yield a difference in kind, but just one in degrees. Hetherington's minimalism thus denies the

significance of the tertiary value problem⁵². That is a fairly unpalatable upshot for Hetherington's gradualism, as it will impose a heavy burden of proof upon gradualists to prove why, *contra* the majority of epistemologists, knowledge does not possess any distinctive epistemic value compared with mere true belief.

One might argue that gradualists should not take the tertiary value problem seriously, as it seems to rest on theoretical assumptions that they deny. That is, absolutism assumes that the conceptual difference between knowledge and true beliefs is one of kind rather than one of degree. Gradualism, at least Hetherington's first version of gradualism, denies this. Knowledge's being gradational implies that we can put knowledge on a kind of continuum connecting to lesser epistemic standings that fall short of knowledge. Correspondingly, we can also put knowledge on a continuum of value with true beliefs that fall short of knowledge. The difference between knowledge's value and that of true beliefs is just a matter of degree. If that is the case, then this can get Hetherington off the hook.

I do not think Hetherington's minimalism can be defended by this response. First, this response does not counterbalance the burden of proof that Hetherington's gradualism has to shoulder for rejecting justificationism (and, in this case, for rejecting the appropriateness of the tertiary value problem of knowledge). Hetherington still has to demonstrate why the conceptual (and thus axiological) difference between knowledge and lesser epistemic standings is just one of degree rather than one of kind. As aforementioned, it

⁵² For literature acknowledging the appropriateness of the tertiary value problem of knowledge, see Greco (2009, 2010); Pritchard, Alan, & Adrian (2010); Ryan (2017); *etc.* For a challenge, see Simion & Kelp (2016). Simion and Kelp call into question a strong claim that knowledge is distinctively more valuable than mere true belief in the way that Pritchard puts. However, they endorse a weaker claim that knowledge enjoys a *weak superiority* over true belief. That is, some amount of knowledge enjoys more value *of a certain kind* than *any amount of* true belief. Notice that, it is also not obvious whether Hetherington's minimalism meets this weaker claim. Within his minimalist picture, it is unclear *what kind* of distinctive value 'good knowledge' (well-justified ones) enjoys over 'minimal knowledge' (mere true beliefs). Instead, it seems that Hetherington was 'putting knowledge on a kind of continuum of value', which is exactly what Simion and Kelp argue against.

is unwise for gradualists to develop gradualism with this kind of unnecessary extra burden (we will see why it is *unnecessary* soon).

Second, more fundamentally, it is unclear why gradualism has to deny that the conceptual difference between knowledge and lesser epistemic standings is one of kind. Knowledge's being gradational does not entail that it is not a distinct *kind* of epistemic standing in contrast with true beliefs. 'Green' is a gradational concept; it can be put on a continuum of spectrum connecting with other colours such as blue and yellow. But this does not affect the fact that green is a distinct *kind* of colour. The conceptual difference between green and yellow is more like a matter of kind rather than merely one of degree. It will be clearer as my own gradualist proposal unfolds (see Chapter 4) that gradualism just needs to deny that there is a clear cut-off point distinguishing knowledge from lesser epistemic standings—but it does not follow that knowledge is thus the same kind of epistemic standing with true beliefs. Concept A's being a distinct kind of concept from concept B does not rely on the existence of a clear cut-off point between A and B. There is no clear cut-off point between 'jumping' and 'flying' (how far/high should one jump so that one can be counted as flying?), and both jumping and flying can be understood as gradable. Nevertheless, flying is a distinct kind of action in contrast with jumping (we do not see a plane as something that can *jump* very high and far). Moreover, flying enjoys its distinctive value that jumping does not have. Therefore, it is not the case that if a concept A is gradable, and A is evaluated along the same continuum of value with concept B, then A and B do not have distinct conceptual differences and their differences in value cannot be one of kind.

So why can a gradualist proposal differing from Hetherington's avoid being plagued by the tertiary value problem? As we have noted at the end of section 1.2, gradualism does not have to be based on minimalism. It is just Hetherington's specific commitment to external absolutism that leads him to choose to combine his gradualism with minimalism. However, if we discard

external absolutism by rejecting the idea that there is a threshold for knowledge, then this will not only lead us to a more thorough version of gradualism, but also enable us to preserve justificationalism and the appropriateness of the tertiary value problem. The idea is, there is no cut-off point distinguishing knowledge from lesser epistemic standings, nevertheless, as shown before, this does not entail that knowledge cannot be distinctively more valuable than true beliefs.

In summary, Hetherington's anti-justificationism fails to offer a solid ground for epistemic gradualism. At least, a version of epistemic gradualism combined with anti-justificationism would give rise to resistance from the mainstream of epistemologists who take justification as a necessary condition for knowledge. Moreover, it would encounter difficulties in explaining the distinctive epistemic value that knowledge enjoys over mere true belief. This is a significant defect of Hetherington's gradualist theory, since just as BonJour points out, any acceptable epistemological theory must make sense of the important value of knowledge (see BonJour 2010).

3.2: Threshold Absolutism

Hetherington insists that there has to be an absolute cut-off point between knowing and not knowing. That leads him to adopt a form of threshold absolutism (in Hetherington's own terminology, 'external absolutism'), which maintains that the threshold between knowledge and anything that falls short of knowledge is absolute. This absolutist view can be further divided into two sub-stances that Hetherington is committed to:

[Trans-Threshold Absolutism]

There is one and has to be one correct sufficient and necessary condition for constituting knowledge.

[Intra-Threshold Absolutism]

There is a clear and absolute cut-off point between satisfying and failing to satisfy the sufficient and necessary condition for constituting knowledge.

Hetherington endorses trans-threshold absolutism in the sense that he defines 'true belief' as the sufficient and necessary condition for knowledge—on one hand, to be qualified as knowledge, an instance of belief has to be true; on the other hand, once an instance of belief turns out to be true, then it suffices to be counted as (at least the minimal) knowledge. He also accepts the intra-threshold absolutism, in that the cut-off point between a belief's being true and being false is ordinarily taken to be clear and absolute. It is not hard to see that Hetherington's threshold absolutist inclination is based on his commitment to his minimalist account of knowledge. However, as we have noted in the last section, it is untenable to define mere true belief as the minimal knowledge. As a result, both Hetherington's trans-threshold and intra-threshold absolutisms would fail: Firstly, given that mere true belief is not a plausible sufficient and necessary condition for knowledge, it is unclear whether there is an absolute sufficient and necessary condition that can distinguish knowledge from non-knowledge in a trans-threshold manner. Secondly, given that justification should also be taken into considerations when judging whether a true belief suffice to constitute knowledge or not, and that Hetherington admits that it is almost impossible to locate a cut-off point between 'sufficient justification' and 'insufficient justification', knowledge will not be able to be absolute in the intra-threshold sense as well.

But is there any other possible approach to defend the two kinds of threshold absolutism that are not based on Hetherington's minimalism? In particular, can the two kinds of threshold absolutism be independent of each other, rather than being as closely related as Hetherington presents them to be? Can they both, independently, indicate that the threshold of propositional knowledge is absolute, respectively?

When claiming a notion to be absolute, one usually means that the notion is absolute in both trans-threshold and intra-threshold sense. Nevertheless, there are exceptions. For instance, 'bilingual' is ordinarily taken to be an ungradable term, a person can either be bilingual or not. That is because, the sufficient and necessary condition for being bilingual is clear—being able to speak two languages. Thus, the threshold between being bilingual and not being bilingual is absolute in a trans-threshold sense. However, the sufficient and necessary condition *per se*, viz, 'being able to speak two languages', is gradational. There is no such a clear and absolute cut-off point that can distinguish 'being able to speak two languages' and 'being unable to speak two languages'. How fluently should one speak that language? Is there a minimal requirement for the size of one's vocabulary? How many grammatical mistakes are tolerable? Answers to these questions are far from clear. Therefore, 'being bilingual' cannot be seen as an intra-threshold absolute term.

The moral that we can draw from the above analysis is that a complete objection to threshold absolutism ought to reject both trans-threshold and intra-threshold absolutism. So apart from that fact that Hetherington's threshold absolutism is based on his problematic minimalist account of knowledge, is there any other reason suggesting that we shall reject threshold absolutism?

Here are some reasons that I take to be important. Firstly, by endorsing the threshold absolutism while only arguing for the beyond-threshold absolutism, Hetherington misjudges the central dispute in the absolutism/gradualism debate. According to Hetherington, what we ought to strive for is just the beyond-threshold gradualism. If this is an anti-orthodoxy attempt of epistemological interest as Hetherington alleges, then it implies that most epistemologists are committed to the beyond-threshold absolutism so defined correspondingly:

[Beyond-Threshold Epistemic Absolutism]

Propositional knowledge is absolute, in the sense that even beyond the threshold between knowledge and not-knowledge, knowledge cannot be better or worse.

However, that is not true. Whereas we have noted in the first chapter that most epistemologists would be happy to endorse that one's knowledge can be improved in the sense that the epistemic support for this instance of knowledge can be strengthened, Hetherington's beyond-threshold gradualist proposal appears to be attacking a straw man. He claims that there is a dilemma for epistemic absolutists as follows:

- '1. Knowledge is absolute, non-gradational. That is, knowledge that p cannot be better or worse as knowledge that p.
2. Each part of one's justification for a belief which is knowledge is contributing somehow to the belief's being knowledge, to how it is the knowledge it is.' (Hetherington 2001: 21)

Hetherington alleges that 1 is incompatible with 2, and thus in order to maintain 1, absolutists have to deny 2, which is an unpalatable result for absolutism as the justification can certainly be better or worse:

'That is, in my previous knowledge-gradualism, the degree or quality of a piece of knowledge was a function of what constituted the knowledge 'from within'—the strength of its justificatory component.' (Hetherington 2005:138)

However, according to our analyses in the first chapter, absolutists can grant 2 and argue that 1 is consistent with 2. The central dispute on the table is not whether one should accept 2 or not, but whether 1 is compatible with 2. By maintaining that 1 and 2 are consistent, absolutists can adopt the equivocal stance. They grant that knowledge can be improved in terms of the quality of justification, but they do not take this kind of improvement to be of

epistemological interest. The fact that knowledge can be 'improved' in this sense does not suffice to indicate that knowledge is gradable. Take the famous analogy of 'pregnant' for instance, Dretske asserts that knowing a fact is like being pregnant. It is not completely unreasonable to argue that one's pregnancy can be seen as improved, *in a certain sense*, if she has healthily and safely come to the late trimester of her pregnancy. However, the length for which one has been pregnant does not affect whether one is pregnant or not. 'Being pregnant' is still a yes-or-no affair. Likewise, it is undeniable that for a student who gets his bachelor degree with a 2:2 and gets his master degree with a distinction, his second graduation, *in a certain sense*, is a *better* graduation. It can even be claimed that the student graduates *in a better way* from his master programme. Nevertheless, 'graduates' is still an ungradable term, an absolute yes-or-no affair in the sense that one can either graduate or fail to graduate. Epistemic absolutists can argue that any concept, including 'knows' for sure, can be improved or be better *in a certain sense*, but that *sense* might just be an insignificant one, which is not of enough epistemological interest to make the concept a gradable one. Hetherington neglects the possibility of this equivocal stance, and thus leads his gradualism to a wrong direction. Focusing merely on the beyond-threshold gradualism also mistakenly diagnoses the reason why the orthodoxy of epistemology does not accept gradualism. The key is not to show that propositional knowledge can be improved *in a certain sense*, but to show *that sense* to be more epistemologically significant. In other words, the key is to reject the equivocal stance.

To be more specific, advocates of the equivocal stance might argue that most ungradable terms can be evaluated among at least two orthogonal scales—one is *essential* (intrinsic, genuine, central, authentic), and another is *extra* (incident, additional, peripheral, accessory). For example, 'defeats' is an ungradable term. 'Whether a team's final score is higher than that of its opponent' can be the essential scale for a team's 'defeating' its opponent, while 'by how many scores the team wins' can be the extra scale. Compared

with a team that wins by 7 scores, another team that wins by 17 scores can be seen as having a 'bigger victory'. However, this cannot affect the fact that both teams *defeat* their opponents, and that 'defeats' is an absolute yes-or-no affair (especially in competitions like basketball games or football games). Analogously, the quality of knowledge can also be evaluated among two orthogonal scales—stronger epistemic support can improve a piece of knowledge in the *extra* sense, but not the *essential* sense. The strength of the epistemic support, once reaches the threshold, would fall into the extra scale and has nothing to do with knowledge's gradability. Just like a larger advantage in scores can make a victory a bigger (and thus 'better') victory, but cannot make the concept 'defeats' *per se* come in degrees. One's epistemic state of knowing-that-*p* can be evaluated as 'becoming better', but only among an extra scale, a scale that does not need to be taken into consideration when answering the 'whether-or-not' question (e.g., 'whether she knows that *p* or not'). Hence, the connection between the absoluteness of knowledge's threshold and the absoluteness of the concept of knowledge is not as close as Hetherington takes it to be.

The second reason why gradualists had better not endorse threshold absolutism is fairly straightforward—that is the threshold absolutist proposal can hardly succeed. Trans-threshold absolutism aims to figure out an absolute sufficient and necessary condition for knowledge, *viz.*, a correct analysis of knowledge. However, the post-Gettier epistemology has encountered continuous frustration in pursuing this goal. Various new analyses of knowledge have been given in response to Gettier's counterexamples (see Gettier 1963) to the classic JTB analysis. At the same time, various new counterexamples continue to emerge in response to those putative correct analyses (see Shope 1983 for a detailed survey). Zagzebski (1994) and Craig (1999) even suggest that the post-Gettier project is hopeless at providing a correct analysis of knowledge. Timothy Williamson's influential book *Knowledge and Its Limits* (2000) and numerous sequent knowledge-first studies further argue that the analysis project should be

replaced by a new proposal 'putting knowledge first'—instead of a notion to be analysed, 'knowledge' should be taken to be the most basic concept prior to any other epistemic concepts. I am sympathetic to these anti-analysis fashion in contemporary epistemology, and I do not think that a necessary and sufficient condition is necessary for the epistemic gradualist proposal. Even if it is possible to reasonably defend a putative analysis of knowledge in a non-circular and non-arbitrary way, basing gradualism on such an analysis is not an ideal strategy, given that there appears to be little consensus on the proper analysis of knowledge by now. A gradualist theory based on a specific putative analysis of knowledge will inevitably encounter resistance from people who disagree with the analysis. Endorsing trans-threshold absolutism will invite unnecessary difficulties for gradualism.

As for intra-threshold absolutism, it is also unclear how to locate an absolute cut-off point between satisfying the sufficient and necessary condition or not. It is unclear whether there is such a condition; even if there was, it is also unclear whether there is such a cut-off point. Absolutists might argue that the cut-off point locates at the boundary between true and untrue, as whatsoever, truth has to be a necessary condition for knowledge given that 'knows' is a factive term. No matter what other factors constituting the sufficient and necessary condition are, we can draw the boundary between knowledge and non-knowledge among the clear and absolute boundary between true and untrue. Therefore, due to the factivity of 'knows', the cut-off point distinguishing knowing from not knowing just locates at the boundary between truth and falsity. However, a factive term does not necessarily have an absolute boundary, and thus can also be gradable. For example, 'understands', at least in the sense that most epistemologists are interested in, is also a factive term. If one understands that p , then p must be true. It is inconsistent to say something like 'I understand that you want to save her life, but in fact, you do not want to save her life'. Notwithstanding, 'understanding' is ordinarily taken to be a gradable concept. Understanding is not merely a yes-or-not affair, rather, it appears to be a matter of degree. One's

understanding can be improved (in an essential sense), and one can also be claimed to understand a fact better than another person does. Factivity thus fails to lend strong support to nominate the cut-off point for the threshold of knowledge. Neither trans-threshold nor intra-threshold absolutism is beneficial to gradualism.

4. Criticism of Hetherington's Second Gradualist Theory

Hetherington's second design for gradualism turns essentially on the following argument:

[Argument from Equivalence]

(E1) Having knowledge that p entails having how-knowledge that p , as knowledge that p is the minimal aspect of how-knowledge that p .

(E2) Having how-knowledge that p entails having knowledge that p , as knowledge that p is the minimal aspect of how-knowledge that p .

(E3) From E1 and E2, having knowledge that p is the necessary and sufficient condition of having how-knowledge that p .

(E4) From E3, having knowledge that p is equivalent to having how-knowledge that p .

(E5) How-knowledge that p is gradable.

(E6) Therefore, knowledge that p is gradable.

This argument is untenable, because: first, its first two premises are based on a mistaken epistemological dogma; second, the schema of this argument is invalid. Correspondingly, there are at least two ways to object this argument.

The first objection focuses on E1 and E2. When proving E1 and E2, Hetherington argues that:

‘Manifestly, that *p* obtains is one of those aspects of how it is that *p*: it is a *minimal* aspect’ (Hetherington 2005:131)

That is to say, in Hetherington’s understanding, knowing that *p* obtains (*i.e.*, that *p* is true) is equivalent to knowing that *p*, so that knowing that *p* can constitute the minimal aspect of knowing how that *p*. E1 and E2 cannot be smoothly entailed if there is a conceptual lacuna between ‘knowing that *p*’ and ‘knowing that *p* obtains’. Therefore, both E1 and E2 imply that one’s ‘knowing that *p* is true’ suffices to guarantee one’s ‘knowing that *p*’.

Remember that in Chapter Two, we have discussed the so-called KT-schema, which maintains that ‘knowing that *p*’ is equivalent to ‘knowing that *p* is true’. We have also seen that KT-schema can be disproved by counterexamples where the subject knows *p* to be true, but still fails to *know that p* due to the lack of understanding regarding the content of *p*. Therefore, the failure of KT-schema suggests that ‘knowing that *p* is true’ does not suffice to guarantee one’s ‘knowing that *p*’. Rather, ‘knowing the content of *p*’ seems to be another complementary factor that should be taken into consideration. That is to say, from what we have discussed, ‘to know that *p*’ seemingly means that ‘to know that *p* is true + to know the content of *p*’. To be specific, I advocate a competing account of the object of propositional knowledge:

[T+C Account] To know that *p* is to know that *p* is *true*, plus, to know *p*’s *content*, which requires understanding the meaning of *p*’s content.

The difference between my T+C account and Hetherington’s second gradualism lies in whether ‘understanding the content of *p*’ is necessary for

'knowing that p '. E1 and E2 are premised on the KT-schema which maintains that merely 'knowing that p obtains' is enough for constituting the minimal 'knowing-that- p '. Accordingly, 'understanding the content of p ' is *unnecessary*. As an opposite, my T+C account argues that one's 'understanding the content of p ' is necessary for her 'knowing that p '. Hence if T+C account turns out to be a more complete answer to the object problem of propositional knowledge, then E1 and E2 should be denied as they are based on the untenable KT-schema. Correspondingly, Hetherington's argument from equivalence will be blocked as its essential premises are falsified.

Even if one insists that KT-schema is true, there is still another way to refute the argument from equivalence. The second objection comes in the form of a *reductio* to Hetherington's argument. The argument from equivalence trades on an underpinning sub-argument as follows:

- (1) A minimal aspect of φ entails that φ , and vice versa;
- (2) Hence a minimal aspect of φ is equivalent to φ ;
- (3) Therefore a minimal aspect of φ should have the same property P that φ has.

In Hetherington's original argument, φ refers to how-knowledge that p , of which the minimal aspect is knowledge that p . Given that how-knowledge that p has the property of 'being gradable', as its minimal aspect, knowledge that p should also be gradable. However, if the schema stated before is valid, then we can prove many unreasonable conclusions. For example:

- (E1*) 'Being dual-purpose' entails 'being versatile,' as being dual-purpose is the minimal aspect of being versatile.

- (E2*) 'Being versatile' entails 'being dual-purpose', as being dual-purpose is the minimal aspect of being versatile.
- (E3*) From E1* and E2*, being dual-purpose is the necessary and sufficient condition of being versatile.
- (E4*) From E3*, being dual-purpose is equivalent to being versatile. (E5*) 'Versatile' is a gradable concept.
- (E6*) Therefore, 'dual-purpose' is a gradable concept.

Being versatile means being able to be used for many purposes. Being able to be used for two different purposes, *i.e.*, being dual-purpose is the minimal aspect of it. Versatile is a comparable and gradable concept, it can be felicitously said that 'this product is more versatile than another', which means 'this product can be used for more purposes than another'. However, dual-purpose is an absolute concept, and it does not make sense to say that 'this product is more dual-purpose than another'. Similarly, 'multicultural', 'multiracial' and 'multilingual' are all gradable, nevertheless, their minimal aspects 'bicultural', 'biracial' and 'bilingual' are ungradable. 'Having two colours' can be the minimal aspect of being colourful, however, colourful is a gradable notion while 'having two colours' is not. 'Can cause one's death' is the minimal aspect of being fatal (at least, a sufficient and necessary condition for being fatal), nevertheless, fatal is gradable ('this is the most fatal disease in the world'; 'this snake is more fatal than that one') but 'can cause one's death' is not. There are many other counterexamples that I cannot exhaust here.

Given these *reductios* and counterexamples, it should be clear that the schema of argument utilised by Hetherington fails to obtain. The argument from equivalence is not a valid argument that can properly derive the gradualist conclusion. Hetherington's second version of gradualism is thus ill-grounded.

5. Concluding Remarks

This chapter critically analysed the most robust and influential epistemic gradualist proposals in the literature, to wit, Hetherington's two versions of gradualism. It has been proved that both the two proposals fail to provide us with a satisfactory account of gradualism. In particular, it seems that by advocating a mixture consisting of threshold absolutism and beyond-threshold gradualism, Hetherington falsely diagnoses the central divergence between gradualism and absolutism. By advocating anti-justificationism, Hetherington imposes an unnecessarily heavy burden of proof on gradualism. These are not only defects of Hetherington's gradualist theories—they also became defects adopted by the current model of the absolutism/gradualism debate. Regrettably, the debate initiated by Hetherington did not attract sufficient attention from the epistemologists. A significant reason for the debate's being underappreciated, I believe, is that the current framework of debate is not constructed properly. In the next chapter, I will remodel the debate between gradualism and absolutism. It will be demonstrated that the reconstructed debate should be focused on the question of whether knowledge-that does have a conceptual threshold.

Chapter Four: Remodel the Debate, Threshold and Spectrum

Abstract This chapter aims to overthrow the equivocal attitude toward the gradualism/absolutism debate and try to prove that gradualism is preferable to absolutism. To achieve this goal, I will take two steps. First, I will revisit and remodel the gradualism/absolutism debate so that the real divergence between the two positions can be clarified. It will be argued that the current debate initiated primarily by Stephen Hetherington fails to capture the essential divergence between gradualism and absolutism, and thus is apt to render the debate trivial and equivocal. By introducing the distinction between a *threshold concept* and a *spectrum concept*, it will be suggested that the new gradualism/absolutism debate should focus on the question whether knowledge should be understood as a threshold concept or a spectrum concept. Absolutism sees knowledge as a threshold concept, while gradualism should argue that knowledge is a spectrum concept. The theoretical merits of this new model of the debate will be shown. Second, I will illustrate the theoretical advantages that gradualism enjoys over absolutism. I will start by showing that gradualism can provide us with a unified account of different types of knowledge, which can solve the so-called *asymmetry problem of knowledge*.

1. The Current Model of the debate

In the first two chapters, we have seen that there is no tenable reason supporting the absolutist dogma—in particular, the mainstay linguistic reason cannot lend sufficient support to absolutism. Nevertheless, the debate around gradualism and absolutism is not so easily finished. There still remains an

impression that we do not need to revise our epistemological orthodoxy that knowledge—that is ungradable. The transformation from the absolutist tradition to the gradualist stance lacks enough motivation. The first chapter summarises this lingering impression as the *equivocal attitude*: an equivocal hybrid stance that accepts both the absolutist orthodoxy and the minimal sense of gradualism (that one’s epistemic support is gradable). People holding this attitude are absolutists—in most cases; but they are gradualists too—just in the minimal sense. To reject this attitude, gradualism needs to complete two tasks: first, to undermine the motivation for holding the absolutist orthodoxy, which we have done in previous chapters; second, to strengthen the motivation for embracing the gradualist view, which is what I aim to accomplish by the remaining parts of this thesis.

So why do current epistemological discussions lack motivations for abandoning the equivocal attitude? One important reason, I think, is that the current model of gradualism/absolutism debate initiated by Hetherington fails to capture the real divergence between the two positions. As we have seen in Chapter 3, a significant trait of Hetherington’s non-absolutist theory is that he only rejects the idea that knowledge—that is absolute in the sense that propositional knowledge cannot be better or worse despite being better or worse justified. However, he is committed to the idea that there is an absolute threshold (*i.e.*, ‘cut-off point’, in his own terminology) for knowledge, which is mere true belief. In summary, Hetherington admits that he is committed to a *local* gradualism that consists of two views:

- ‘(i) Knowledge is to be absolutely distinguished from whatever is not knowledge (whatever is external to knowledge). This is because there is an absolute cut-off point between knowing and not knowing,
- (ii) But *within* the category of knowing, non-absolutism is true. That is because it is possible that some cases of knowledge that *p* are better as knowledge that *p*, or more clearly knowledge that *p*, than other cases of knowledge that *p*.’ (2001: 7)

View (i) is called ‘external absolutism’ by Hetherington, and view (ii) is named ‘internal non-absolutism’. Given Hetherington’s anti-justificationist position, it is not hard to understand why he endorses external absolutism. For Hetherington, the central divergence between gradualism and absolutism is not whether there is an absolute threshold for knowledge, but whether knowledge can be improved *beyond the threshold*. Correspondingly, Hetherington concludes that the answer is certainly ‘yes’—knowledge can be better or worse in terms of the strength of the epistemic support (e.g., the quality of the justification).

Notice that Hetherington’s commitment to external absolutism is reflected in not only his first but also his second version of gradualism. Hetherington (2011) advocates the so-called *knowledge-as-ability hypothesis* arguing that knowledge is an *ability* to manifest various accurate representations of *p*. An ability can be applied better or worse, notwithstanding, Hetherington still emphasises that his gradualism is compatible with there being a threshold for having the ability:

‘In general, abilities are gradational This is consistent with there being an absolute cut-off point, either precise or not, between having the ability and not having it. (So that is not the sense of knowledge-absolutism I am denying.)’ (Hetherington 2011: 49)

So, in what sense does Hetherington’s second version of gradualism deny absolutism? Hetherington (2011) defines absolutism as follows:

‘Knowledge-absolutism thus implies that there cannot be two instances of knowledge that *p*, one of which is somehow a better or higher grade of knowledge that *p* than is the other. So, in particular, no instance of knowledge that *p* is ever improvable purely as knowledge that *p*.’ (ibid: 6)

Therefore, here, Hetherington still discusses absolutism in the ‘internal’ sense, not only because he argues that his gradualism is compatible with

there being a threshold for knowledge (as we have noted), but also because he maintains that absolutism is a view that the quality of an instance of knowledge that *p* (qua knowledge that *p*) cannot be improved by stronger epistemic support (as stated in the quoted passage above), which coincides with his definition of 'internal absolutism' in his early works.

So far, we have seen a brief sketch of Hetherington's epistemic gradualism that has two significant features: 1) being based on anti-justificationism; 2) being committed to external absolutism. Although the gradability problem of knowledge—that is theoretically interesting and significant, Hetherington's gradualism did not receive sufficient echoes as one might expect it did. This is, I think, largely due to its aforesaid two features, as the two features render Hetherington's gradualism both too controversial and too uncontroversial:

On the one hand, the anti-justificationist basis for Hetherington's gradualism severely undermines its appeal, given that the majority of epistemologists still see justification as necessary for knowledge. It is not hard to conceive how Hetherington's courageous defence for anti-justificationism would invite heavy criticism. For example, the true belief account of knowledge is accused of violating the widely-shared intuition that knowledge is incompatible with mere epistemic luck (see Pritchard 2007; Sosa 2007; Greco 2003). The idea that propositional knowledge is gradable is already novel, so when this idea is based on an even more novel and controversial ground, it is understandable that people are not willing to pay serious attention to the talk of gradualism.

On the other hand, as we have noted in the first chapter, many commentators also notice that Hetherington's commitment to external absolutism seems to obscure the distinction between gradualism and absolutism. *Contra* what Hetherington alleges, many epistemologists (e.g., Feldman 2002; Brower 2004; Jaster 2013; Leite 2006; as introduced in Chapter 1) should already be glad to consent that, at least in a loose sense,

there can be better knowledge in terms of better justification—hence Hetherington’s local gradualism is not as unorthodox as he presents it to be.

Moreover, some even hold that the gradability of justificatory strength/probability/fallibility of knowledge is irrelevant to whether knowledge *per se* is gradable. For example:

‘There is no reason, given fallibilism at least, to deny that you can know with more or less probability, but this doesn’t imply that knowledge is gradable in anything like the way that, say, height is or affection is.’ (Fantl & McGrath 2009: 23-24)

Therefore, it seems that most of the absolutists grant that the quality of one’s knowledge can be improved by virtue of better justification—nonetheless, this does not imply that the concept of knowledge is gradable. Hetherington seems to misdiagnose the central divergence of the debate between gradualism and absolutism, and thus faces the danger of trivialising the debate. A direct upshot of Hetherington’s local gradualism is an equivocal attitude towards the gradualism/absolutism debate. The idea is, knowledge is gradable in one sense (*viz*, internal gradualism), while ungradable in another sense (*viz*, external absolutism). Thus, it is unclear whether, overall, knowledge is gradable or not given these two distinct senses. Call this *the two-sense problem*. Hetherington’s model of the debate is silent on *which sense* is more essential for our understanding of the concept of knowledge. Just like ‘jump’—it is ungradable in the sense that there is an absolute threshold for judging whether one jumps or not (*e.g.*, both feet being off the ground); however, ‘jump’ seems to be gradable in the sense that one can jump higher or further than another person does. So is ‘jump’ gradable? The answer might be equivocal. Hetherington reads knowledge-that in a very similar way: on the one hand, knowledge-that is ungradable in the sense that there is an absolute threshold (*i.e.*, true belief) distinguishing knowledge from ignorance; on the other hand, knowledge-that is gradable in the sense that, beyond the threshold, one can have better knowledge if one has better

justification. The two-sense problem thus is apt to result in *equivocalism* regarding the gradualism/absolutism debate—propositional knowledge is simultaneously gradable and ungradable in two different senses respectively. Wherever equivocalism is adopted, there is no genuine debate, as nothing is actually irreconcilable in the putative ‘debate’.

Nevertheless, the gradualism/absolutism debate initiated by Hetherington is still worth revisiting. After all, Hetherington’s investigation is successful in revealing that:

Point 1: The absolutist orthodoxy seems to be ill-grounded. Thus, its opposite, namely gradualism, is worthy of more careful considerations.

Point 2: There are notable resemblances between knowledge-that and other epistemological notions that are paradigmatically gradable, *e.g.*, understanding and knowledge-how. This seems to indicate that knowledge-that might be homologous with those gradable notions.

Point 3: For epistemic absolutists (in particular, fallibilist external absolutists) who insist that knowledge has a threshold, the *boundary problem of knowledge* is a troublesome issue.

All the three points provide us with reasonable motivations to reconsider the debate between absolutism and gradualism. The first two points indicate that gradualism might not be such a crazy view as it appears to be. Absolutism primarily gains support from linguistic evidence, but we have noted in Chapter Two that there are also linguistic data supporting gradable/comparable uses of ‘knows-that’ construction. Apart from counterevidence provided by Hetherington (2001) and Dutant (2007) based on the English language, we can find more linguistic support for gradualism from other languages. For example, in Chinese, it is felicitous to say

something like ‘我父亲比我更知道生活不易’⁵³, where ‘知道’ (the direct translation of ‘knows that’ in Chinese) is used as a gradable term. In addition, it is questionable whether a concept’s linguistic uses can reveal the conceptual nature of an epistemological concept. For example, Hazlett (2010) argues that linguistic evidence suggests that knowledge should be seen as non-factive, given the notable amount of non-factive uses of ‘knows-that’ in the English language⁵⁴. Nevertheless, it is almost an epistemological consensus that knowledge implies truth, and hence knowledge, by its nature, is factive. Hazlett thus concludes that there is a gap between an epistemological notion’s linguistic uses and its conceptual nature. In summary, the linguistic motivation for absolutism is untenable. In contrast, given point 2, it is intriguing for us to give gradualism a more serious consideration. Perhaps embracing gradualism can provide us with a unified account of knowledge that eliminates the somewhat odd asymmetry between different types of knowledge—why is propositional knowledge is ungradable while knowledge-how and knowledge-wh are widely accepted to be gradable? This asymmetry is especially knotty for epistemologists (e.g., intellectualists) who believe that different types of knowledge can be reduced to one single type of knowledge. Gradualism might give this issue a straightforward resolution: the asymmetry just does not exist, because knowledge-that is gradable as well.

Point 3 offers a ground for doubt on the very existence of knowledge’s threshold. The idea is, if neither infallibilism nor Hetherington’s anti-justificationism is an appropriate solution to the boundary problem, why do not consider the possibility that knowledge-that does *not* have a threshold?

⁵³ It means ‘my father knows better that life is hard than I do’.

⁵⁴ Bac & Irmak (2011) also illustrates a similar point on the basis of the Turkish language.

Hetherington, as well as many other epistemologists⁵⁵, seems to accept the existence of knowledge's threshold as an undeniable fact—the problem is not whether this threshold exists, but how we can (even if only approximately) locate that threshold. There seems to be a close connection between the idea that *knowledge has a threshold* and that *knowledge is fallible*:

'[K]nowledge's having a justificatory boundary is a vital part of all traditional fallibilist conceptions of what knowledge is.' (Hetherington 2006: 47)

Hetherington argues that fallibilism naturally invites the question: how fallible is a piece of knowledge allowed to be? Fallibilism thus perforce gives rise to the boundary problem. Hetherington also realises that it is extremely difficult for fallibilists to locate a precise threshold of knowledge. He also admits that our ignorance of the location of knowledge's threshold can significantly undermine our knowledge that knowledge has a threshold. Nevertheless, he still insists that this does not affect the fact that the threshold exists⁵⁶. At least, there is one achievable way of solving the boundary problem that fallibilists can take, that is, to 'narrow the area or span' within which the threshold lies. This attitude is also shared by Michael Hannon:

'While an adequate answer to the threshold problem shouldn't demand more precision than is achievable (or necessary), fallibilists must provide a reasonable degree of approximation. They have not met this challenge, and so I take the threshold problem seriously.' (2017: 608)

Therefore, Hetherington's *local gradualism* is thus expected to find an approximation of where knowledge's threshold is. It never questions the very existence of such a threshold. However, as we have seen before, this would invite the two-sense problem and trivialise the debate between gradualism and absolutism. It seems that *external gradualism*, rather than internal

⁵⁵ For example, BonJour (2003, 2010), Fantl & McGrath (2009), Bovell (2012), Hannon (2017). Rothschild & Spectre (2018) also defend this threshold view of knowledge on the basis of a sorites objection.

⁵⁶ Similar ideas are echoed by Hannon (2017) and Bovell (2012).

gradualism, is what genuinely controversial. So, can we remodel the gradualism/absolutism debate by appealing to a *global gradualism* instead, which focuses on not whether knowledge can be improved beyond the threshold, but whether knowledge has a threshold? In what follows, I will develop such a new version of gradualism which interprets knowledge as a *spectrum concept* instead of a *threshold concept*. By doing so, we can usher a more promising way toward remodelling the absolutism/gradualism debate.

2. Threshold vs. Spectrum

2.1: Threshold Concepts

What does it mean by saying that a concept has a threshold? Ordinarily, 'threshold' is understood as a point on a scale at which something starts to happen or change. If something is taken to be of a threshold, then usually this concept can be evaluated or measured along a scale in a certain sense. For example, when we say that someone has a high *threshold for pain*, what is evaluated is the strength of pain that one can endure—when the strength reaches the threshold, one would be unable to endure the pain. When we say that 60% of the vote is the *threshold for approval* of a proposal, what is measured is the percentage of the vote that a proposal can gain—when the percentage reaches the threshold, the proposal can be approved. Thus a concept that has a threshold ought to be evaluable along a certain scale, where its threshold is a point where anything starts counting as an instance of that concept. Accordingly, let us define:

[Threshold Concept]

A concept C is a threshold concept, *iff*, at the relevant scale S on which C is evaluated, there is a cut-off point T that distinguishes instances of C from everything that falls short of C.

A threshold concept might be evaluated on multiple scales. For example, if knowledge is taken to be a threshold concept, then relevant scales can include the scale of belief, the scale of justification, and so on. I am also glad to admit that the trade-offs between these different scales can settle an *overall threshold*, and this overall threshold typically comes in the form of necessary and sufficient conditions. For the sake of simplicity, the definition of threshold concept above only employs a single scale S—however, the scale S here can also be understood in the form of a conjunction that combines multiple scales together. External absolutists read knowledge as a threshold concept. In contrast, I will argue that gradualism should reject this reading. Before we move forward to applying the notion of ‘threshold concept’ to knowledge, some caveats about this notion are in order:

First, albeit evaluable/measurable on a certain scale, a threshold concept is not thus perforce gradable. Quite the opposite, numerous threshold concepts are taken to be ungradable, *e.g.*, ‘empty’, ‘extinct’, ‘completed’, and ‘sold out’. For these concepts, the threshold is located at the endpoint of the scale and thereby renders these concepts absolute—either reaches the threshold or not (for example, either empty or not empty), it is a yes-or-no affair. That is also why the threshold reading of knowledge matches external *absolutism* rather than gradualism.

Second, a term can be ungradable without thus being a threshold concept, if there is no un-farfetched corresponding *scale* along which the term can be evaluated. For example, ‘wooden’, ‘foreign’, and ‘the third’.

Third, when saying that there is a ‘cut-off point’, I do not deny the fact that a threshold concept can still have its borderline cases. In other words, ‘threshold’ is compatible with ‘vagueness’, given that a vague concept is ordinarily understood as a concept that has borderline cases. There are at least three ways to understand how a threshold concept can be vague. The first way appeals to the *practical* difficulty in telling whether the cut-off point is

reached. Whether a cut-off point exists is one thing, while how to judge whether the cut-off point is reached is another thing. A threshold concept C's threshold provides a criterion for us to determine (in principle) what is C and what is not. Nevertheless, there can be some sub-criteria for determining (in practice) whether the threshold is met, and those sub-criteria can be much vaguer. By way of illustration, in the UK, a student's essay 'fails' *iff* its mark is below 40. In that case, 'fails' is a threshold concept of which the cut-off point falls at score 40. However, whether an essay should be given a score below 40 could be a headache for teachers. When it comes to the evaluation of an essay's quality, the rubric can be highly vague.

The second way is to appeal to the *cognitive difficulty* in detecting the threshold by adopting *epistemicism* (see Williamson 1994, 2007; Sorensen 1988; Horwich 2000; Hu 2015; *etc.*). Epistemicism of vagueness argues that there is a sharp boundary between the extension and the anti-extension of a vague predicate. For example, there is a cut-off point from which a person can be recognised as 'bald'. It is just because of our cognitive inability that it is impossible for us to know where the threshold is—nevertheless, this does not affect the fact that the threshold exists. Accordingly, threshold is compatible with vagueness, where vagueness is interpreted as ignorance of the location of the cut-off point. Hetherington's attitude towards the boundary problem of knowledge is very akin to an epistemicist reading: there is a sharp boundary for knowledge, and our ignorance of where the boundary is located just undermines our *knowledge* that there is a threshold, but it does not undermine the *fact* that there is a threshold (see Hetherington 2006).

The third way is to appeal to fuzzy logic that admits *degrees of truth*. One can endorse a form of infinite-valued logic in which the truth value of a proposition F_φ (φ stands for any potential instance of F) varies from 0 (definite falsehood) to 1 (definite truth) inclusive. The predicate F is vague, as any F_φ whose truth value is a real number n ($0 < n < 1$) that is very close to 1 can be seen as a borderline case of F (see Smith 2008; Hyde 2010; Shapiro 2006; *etc.*).

Correspondingly, F 's threshold can be interpreted as the point where F_ϕ starts being designated to the truth value of 1, *viz.*, the right endpoint of the closed interval $[0, 1]$ ⁵⁷.

Finally, a threshold concept can be context-sensitive in the sense that its threshold might vary from context to context. A male athlete's being 'heavyweight' is a threshold notion. However, the threshold varies in different sports—say, for boxing, the minimum is 200 pounds, while for wrestling, it can be 184 pounds. So when a ring announcer describes an athlete as 'heavyweight', the threshold depends on the context where the statement is asserted.

I hope that I have given a clear enough sketch of the idea of 'threshold concepts'. So what does it mean to see *knowledge* as a threshold concept? Remember that a threshold concept should be able to be evaluated on a corresponding scale. When it comes to the threshold for *belief*, the scale can be the degree of *credence*. But when it comes to the threshold for knowledge, what is/are the scale(s) that we should take into consideration? There can be plural answers to this question. For example, there can be multiple scales in terms of knowledge's various necessary conditions. Most current discussions around the threshold problem of knowledge focus on the scale of justificatory strength, in particular, the degrees of fallibility. Besides, insofar as we take belief to be necessary for knowledge, the scale of *credence* for *belief* can show us another aspect of the threshold problem of knowledge. In this aspect, there are disputes about where the threshold for sufficient credence for yielding knowledge-apt beliefs is. The *credence-one view* argues that beliefs require the maximal confidence or credence 1 (see Clarke 2013, Greco 2015, *etc.*), while the *threshold view* holds the beliefs

⁵⁷ Naturally, one might ask whether there is a sharp boundary for a borderline case's being borderline. That is to ask, are there borderline cases of borderline cases of a vague predicate F ? This query relates to the issue that whether higher-order vagueness exists. An absolutist can accept the idea of higher-order vagueness. Alternatively, she can also argue that higher-order vagueness is just an illusion (see Wright 2010). Interesting as it is, this issue is beyond this chapter's purpose.

only require credence above a threshold that is below degree 1 (see Kaplan 1996, Foley 2009, *etc.*). It is noteworthy that although the credence-one view is seen as an opposite to the threshold view, it does not deny the existence of belief's threshold (in terms of the degrees of credence). The disagreement between the two competing views lies in *where the threshold is*—degree 1 or a certain number that is below 1? Apart from these, various other interpretations of knowledge's threshold can be provided in accordance with different accounts of knowledge. For instance, anti-luck virtue epistemology advocates an analysis of knowledge as follows: 'S knows that p if and only if S's safe true belief that p is the product of her relevant cognitive abilities (such that her safe cognitive success is *to a significant degree creditable* to her cognitive agency)' (Pritchard 2012a: 20; emphasis mine). Accordingly, the threshold for knowledge can be understood as a cut-off point indicating how *creditable* is *creditable enough* for yielding knowledge. This chapter will discuss knowledge's threshold in a general sense, that is, without being limited to any specific account of knowledge⁵⁸.

With these being said, taking knowledge to be a threshold concept means being committed to this thesis:

[K-THRESHOLD]

Knowledge is a threshold concept in the sense that there is a threshold T_k , which distinguishes knowledge from everything that falls short of knowledge.

As we have noted before, many epistemologists who endorse K-THRESHOLD also admit that it is very hard (if not impossible) for us to know where T_k is (*e.g.*, see BonJour 2003; Hetherington 2006; Bovell 2012; Hannon 2017; Foley 2009). Call this attitude *moderate K-THRESHOLD*, which might be a position that most epistemic absolutists adopt. In contrast,

⁵⁸ However, for the sake of convenience, if helpful, readers are welcomed to primarily interpret the threshold mentioned in this chapter in terms of justificatory strength as the majority of current literature does.

there are also voices claiming that we *can know* exactly where the threshold is (e.g., see Hetherington 2001; Goldman 1993; Clarke 2013; Greco 2015). Call this attitude *radical K-THRESHOLD*. In addition, we have noted that a threshold concept can also be context-sensitive, a third attitude defending K-THRESHOLD argues that knowledge has its threshold, while T_k can vary from context to context in accordance with practical factors, such as stakes (e.g., see Fantl & McGrath 2009; Hannon 2017; Hawthorne 2003; Stanley 2005)⁵⁹. Call this attitude *contextualist K-THRESHOLD*. Note that the three attitudes above are not mutually exclusive. Rather, both moderate and radical K-THRESHOLD can overlap with contextualist K-THRESHOLD. Endorsing any version of K-THRESHOLD means being committed to external absolutism.

In the previous chapter, we made a distinction between trans-threshold absolutism and intra-threshold absolutism. K-THRESHOLD reflect the gist of the both types of absolutism. That is, T_k typically comes in the form of an absolute necessary and sufficient condition of knowledge (trans-threshold), and there is a cut-off point for this threshold's being reached (intra-threshold). K-THRESHOLD is a *prima facie* plausible thesis that is deeply entrenched in standard epistemological thinking. K-THRESHOLD's apparent plausibility comes from at least three aspects: first, linguistic data suggest that it is often infelicitous to use 'knows-that' as a comparable construction (see Stanley 2005; Dutant 2007; etc.), in the sense that 'knows' is merely a 'yes or no' affair, where the distinction (threshold) between 'yes' and 'no' ought to be clear. Second, most epistemologists are fallibilists who hold that knowledge is allowed to be fallible *to a certain degree*. While if there is no such a threshold for knowledge, that means we cannot find *the certain degree* to which knowledge is allowed to be fallible, which might significantly undermine the persuasive force of fallibilism (see Hetherington 2006; Hannon 2017). Third, we find knowledge distinctively more valuable and desirable than every

⁵⁹ They advocate a solution to the boundary problem of knowledge called *impurism* – the threshold for sufficient justificatory strength is not only determined by truth-conducive factors, but also partly by one's practical reasoning situation.

epistemic standing falling short of knowledge such as true belief (see Pritchard, 2009a; Bonjour 2010; *etc.*). Correspondingly, it seems to be natural for us to infer that there is a cut-off threshold for knowledge so that knowledge can enjoy its distinctive epistemic value.

Now we have outlined the absolutist side of the remodelled debate by employing K-THRESHOLD. So what about the gradualist side? How can we conceive of a concept of knowledge with no threshold? I suggest that gradualists can resort to the idea that knowledge is a *spectrum concept*.

2.2: Spectrum Concepts

What is the threshold for a colour's counting as red? Where is the boundary distinguishing 'red' from 'yellow', 'blue' or any other colour? It is not only practically difficult to draw this boundary, but also seemingly impossible *in principle* to locate such a threshold. That is because, 'red', as well as other colours, refers to a range of colour across a spectrum, of which the 'boundary' is too gradient to be regarded as a threshold. Admittedly, if one is taken through colour samples that start with blue and become gradually more reddish, there would be some point at which one starts describing them as red. Nevertheless, this only proves that we are able to tell a spectrum concept A from another spectrum concept B, but not that there is a threshold distinguishing A from B. Consider eye tests. In an eye test, as the doctor adjusts the phoropter, the patient who stares at a picture through the lens will be constantly asked whether the vision is clear now, until the patient replies 'yes'. However, this does not imply that there is an absolute threshold for a picture's being 'clear' or 'vague'. When asking for a threshold for 'clear' or 'red', one is committing a category mistake fallacy. A concept's threshold is ordinarily used to identify whether a given instance is a member *of* that concept. However, we do not identify a colour as red by resorting to the so-called 'red's threshold'. Instead, we usually determine whether a colour is an instance of red by (roughly) comparing it with *paradigmatic cases* of red. The

more that colour resembles those paradigmatic cases of 'red', the more likely the colour will be recognised as red. Also, there is no clear threshold for a case's being paradigmatic. Identification of a concept like 'red' is not a matter of 'yes' or 'no', but more like one of degrees.

So is 'cold'/'hot'. A cup of water's being 'ice-cold', 'cold', 'warm', or 'hot' is not a matter that can be determined by a given threshold. These grades of 'warmness' refer to different ranges of temperature on a spectrum of which there are no cut-off beginning points or endpoints. The spectrum of temperature is continuous, gradient, and intrinsically gradable. This spectral nature also applies to many other perceptual concepts, such as 'spicy', 'bright', and 'fragrant'.

Now we have roughly classified a cluster of concepts such as 'red', 'cold', and 'clear', which should be better understood as a matter of spectrum rather than something with a threshold. They should be categorised as a type of concepts that are opposite to a *threshold concept*. Accordingly, let us define:

[Spectrum Concepts]

A concept C is a spectrum concept, *iff*, at the relevant scale S on which C is evaluated, instances of C fall into a range of more paradigmatic or more borderline cases that come in different degrees, while there are *no* cut-off points distinguishing instances of C from anything that falls short of C.

Again, some caveats would be helpful. First, the relationship between spectrum concepts and 'vague concepts'. It is an attractive view for a lot of philosophers that every philosophically interesting concept is somewhat *vague* in this or that sense. So, is every 'vague' concept a spectrum concept? No. Being a spectrum concept is just one way for a concept to be 'vague', but there are still other senses in which a concept could be seen as vague. As we

have noted before, there are some vague threshold concepts. Besides, some concepts that are neither threshold nor spectrum concepts can also be inclined to be vague if they are defined in a functional way. For example, ‘pen’, ‘desk’, and ‘fork’. Where is the boundary between a tree twig and a pen or a fork? Can a large flat rock be counted as a desk? The boundaries are fuzzy. However, they will *not* be spectrum concepts if there are no relevant scales on which those concepts can be evaluated and graded. A spectrum concept requires a gradable scale intrinsic to the concept’s nature. It does not make sense to say something like ‘this desk is more desk than that one’, because no relevant gradable scale entrenched in the concept of ‘desk’ could be found—even though the boundaries of desk’s definition are fuzzy⁶⁰.

Second, are all context-sensitive concepts spectral? No. As aforesaid, there are context-sensitive threshold concepts that are not spectrum concepts. Also, some indexical terms such as ‘I’ and ‘this’ are context-sensitive, but they are neither threshold nor spectrum concepts due to the lack of relevant scales.

Third, is every gradable concept spectral? Given that a spectrum concept has to be gradable, the two categories (*i.e.*, gradable concepts and spectrum concepts) overlap with each other to a very large extent. However, there could be exceptions. That is, a concept might be graded not on a continuous scale but only in a discrete sense. In other words, the concept’s different grades do not constitute a spectrum, but a ‘staircase function’⁶¹. For example, academic degrees. One might see ‘academic degree’ as a gradable notion in that it comes in different levels, *e.g.*, bachelor, master, and PhD. However, for each particular level, there is a threshold for being

⁶⁰ One might also find my definition of ‘spectrum concepts’ somewhat vague. However, this does not mean that the concept ‘spectrum concepts’ is a spectrum concept – because of the lack of relevant gradable scale.

⁶¹ Sosa (2001, 2015) interprets knowledge in a highly similar way. He holds that knowledge can be graded into three levels: animal knowledge, reflective knowledge and knowledge full well. However, one might also argue that, within each level, there is a corresponding threshold for *that level of knowledge*.

awarded such a degree to the effect that anything in between does not count as a *higher degree* (e.g., PhD candidate is not a degree but just a status).

So what does it mean by regarding knowledge as a spectrum concept? Contrary to K-THRESHOLD, the spectrum reading of knowledge holds that in a given relevant scale, there is no such thing serving as the threshold for knowledge. On the contrary, there are better or worse (and of course, more or less paradigmatic) cases of knowledge that come in different degrees which constitute a spectrum. Again, there are various relevant scales that could be taken into consideration when evaluating the quality of a piece of knowledge, e.g., justificatory strength, credence, reliability, the creditability (of one's cognitive success to her cognitive agency)⁶². To put it more clearly, as opposite to K-THRESHOLD, which reflects epistemic absolutism's core thesis, epistemic gradualism should strive to defend the following thesis:

[K-SPECTRUM]

Knowledge is a spectrum concept in the sense that there is *not* a threshold T_k distinguishing knowledge from everything that falls short of knowledge, but only better or worse instances of knowledge that can be graded into different degrees on the relevant scale S .

We have seen that there are three attitudes supporting K-THRESHOLD. Since K-SPECTRUM denies the existence of T_k , it is compatible with neither moderate nor radical K-THRESHOLD. As for the impurism, K-SPECTRUM does not need to deny that knowledge can be context-sensitive, nevertheless, it denies that there are thresholds for knowledge which vary from context to context, as gradualism insists that in no context there exists such thing as T_k . Hence, K-SPECTRUM is opposed to all three versions of K-THRESHOLD.

⁶² Also, for the AAA-model of virtue epistemology, the scales can be about accuracy, aptness, and adroitness.

So how can we distinguish knowledge from what falls short of knowledge if no threshold is available for us to refer to? By endorsing K-SPECTRUM, gradualism can reply that: it is just like how we tell whether an object is red or not. Admittedly, there can be no easy and precise criteria when making judgments like these. In our real-life epistemic practice, judgments about 'whether one knows' are not always easy and precise. Instead, when encountering some borderline cases, our judgments are inclined to be full of uncertainty and hesitation—just like when we see a half-ripe apple. The spectrum view of knowledge can mirror and explain away this hesitation. Nevertheless, it is undeniable that we all have various types of paradigmatic instances⁶³ of knowledge in mind, e.g., logical knowledge and mathematical knowledge of analytic truths (such as '1+1=2', 'P→P'), perceptual knowledge ('I see that I have hands', 'I feel a headache now'), and the like. These cases constitute some typical indicators of paradigmatic instances of knowledge. Also, there are some typical indicators of not-knowledge, such as falsity, disbelief, or mere epistemic luck. It is noteworthy that this also indicates that K-SPECTRUM does not imply that knowledge has no necessary conditions. The spectrum reading of knowledge is compatible with the existence of necessary conditions for knowledge especially when they serve as indicators of paradigmatic not-knowledge. With these indicators of paradigmatic cases of knowledge and not-knowledge in play, we can construct the *frame of reference* for making knowledge ascriptions. The more fine-grained and elaborative the frame of reference is, the better we can accordingly distinguish knowledge from not-knowledge—although the distinction can never be clear-cut or easy given that there are no absolute thresholds.

⁶³ Analogously, consider the 'doneness' of steaks. There are no thresholds for a steak's being 'rare', 'medium rare', 'medium', or 'well done' – one can only (roughly) judge the doneness by how 'pink' or tender the meat is. Restaurants or recipe books might provide you with illustrative pictures showing what steaks with different degrees of doneness would typically look like, which could help you to judge whether my steak is medium rare or medium. Nevertheless, these illustrative pictures are not 'thresholds' for doneness, but just paradigmatic cases for you to refer to.

So is there a threshold for a case of knowledge's being paradigmatic? I think, a consistent gradualist should say 'no'. That is to say, a thorough spectrum account of knowledge ought to endorse higher-order spectrum-ness. Otherwise, K-SPECTRUM is just delaying the inevitable—it just transfers the problem from finding a threshold for knowledge to finding a threshold for 'paradigmatic cases of knowledge'. This does not solve the problem, but just postpones it. One may worry that: does this mean that you are committed to paradigmatic cases of paradigmatic cases of knowledge? And therewith, paradigmatic cases of paradigmatic cases of paradigmatic cases of knowledge, and so forth? Does this invite infinite regress? My reply is: there is nothing wrong with this infinite iteration of higher-order spectrum-ness. A spectrum, by its nature, is the result of an infinite series of gradual changes and the infinite iteration of threshold-less-ness. There is no threshold for red, just like there are no thresholds for crimson, scarlet, laky or claret. One should not worry that a spectrum concept has no higher-order threshold, just like one should not worry that a circle has no straight sides. One should not try to analyse a curve line into straight lines, likewise, one should not expect to analyse a spectrum concept into threshold concepts.

Qua a core thesis that epistemic gradualism should defend, K-SPECTRUM perfectly captures motivations for embracing the gradualist reading of knowledge. As introduced before, gradualism can be motivated by the following discoveries: first, the absolutist orthodoxy is ill-grounded; second, there are significant resemblances between knowledge-that and other paradigmatic gradable epistemic notions. Furthermore, the intractability of the boundary problem of knowledge also renders epistemic gradualism an apparently more intriguing solution. These motivations for epistemic gradualism all lend support to K-SPECTRUM. If knowledge is a spectrum concept, then it is no wonder that there are linguistic data using 'knows-that' as a gradable phrase. Likewise, it can also explain why knowledge-that notably resembles paradigmatic gradable epistemic notions such as knowledge-how and understanding—they are homologous in the sense that

they are all spectrum concepts. Finally, insofar as K-THRESHOLD is accepted, the boundary problem ought to be thereby resolved because it makes no sense to search for a spectrum concept's threshold—spectrum concepts are anti-threshold by their nature.

With K-THRESHOLD and K-SPECTRUM in play, we can now reconstruct the debate between epistemic absolutism and epistemic gradualism by basing it on these two contradictory theses. That means, the central disagreement between absolutism and gradualism should be understood as the dispute the threshold reading of knowledge and the spectrum reading of knowledge. How does the gradualism/absolutism debate benefit from this reconstructed model?

3. Remodelling the Debate

The extant model of the debate between gradualism and absolutism rests essentially on the distinction between *external gradualism/absolutism* and *internal gradualism/absolutism* (*a la* Hetherington). Hetherington argues that external absolutism is unquestionable, and thus we should focus on the comparison between *internal gradualism* and *internal absolutism*. We have noted that this model of the debate has two defects. First, the internal sense of gradualism is not as controversial as Hetherington presents it to be—instead, most absolutists would be glad to admit that knowledge can be seen as gradable in terms of its justificatory strength. Nevertheless, they would still insist that they are epistemic absolutists rather than gradualists, because knowledge is ungradable in another more important sense—it has an absolute threshold. This leads to the second defect of the extant model of the debate between gradualism and absolutism—it gives rise to the *two-sense problem*. Second, within Hetherington's framework, both internal absolutism and (his) internal gradualism are unappealing. Few epistemologists would actually insist that one's knowledge cannot be improved despite better justification—there is apparently no reason for one to defend this

unreasonable stand. However, the internal gradualist reading of knowledge, which is supposed to be relatively more plausible, is based on a fairly unpopular anti-justificationist ground by Hetherington. Consequently, although *external absolutism* is more widely accepted and relatively more debatable such that it is more suitable for constituting one side of the absolutism/gradualism debate, it is hard to find its opposite within Hetherington's framework. That is because, Hetherington did not spell out what an attractive *external gradualist* reading of knowledge is like, because he grants external absolutism; but on the other hand, his *internal gradualism* is too weak to constitute a balanced debate.

By reconstructing the debate on the basis of two irreconcilable theses, to wit, K-THRESHOLD and K-SPECTRUM, the new model of the debate enjoys advantages over the current one in at least three aspects:

- (1) The divergence between gradualism and absolutism is characterised more explicitly. The central disagreement between K-THRESHOLD and K-SPECTRUM is whether there is a threshold for knowledge. Absolutism claims 'yes', gradualism argues 'no'—the divergence is clear and certain. There is no space for equivocalism, as a concept cannot be both thresholded and spectral. Hence it avoids the two-sense problem of the extant debate, and can thus prevent the equivocal attitude.

- (2) The divergence between gradualism and absolutism is captured more accurately. Drawing on Hetherington's own definition of 'external *absolutism*', this new model bases absolutism on its commitment to the existence of a threshold for knowledge. Moreover, for different scales to evaluate knowledge, there is a notable amount of philosophers who advocate the existence of T_k on relevant scales, which prevents gradualists from tilting at windmills. By way of illustration, for the scale of belief, we have

philosophers who endorse credence-one or threshold-view of belief. For the scale of truth, we have numerous bivalentists and epistemicists. For the scale of justification, we have infallibilists (see Williamson 2000; Dutant 2016; Dodd 2011; Littlejohn 2008; Davies 2018; *etc.*), anti-justificationists (*e.g.*, Hetherington and Goldman) and a large number of aforementioned proponents of *moderate K-THRESHOLD* who might be seen as epistemicists on the boundary problem of knowledge.⁶⁴ Compared with Hetherington's model of the debate, of which the 'internal absolutism' side has no genuine proponents, the remodelled absolutism can avoid becoming a straw man.

- (3) The remodelled debate is more balanced as both sides of the debate become more attractive. On the one hand, the remodelled gradualism has more appeal as it does not need to be based on the counterintuitive anti-justificationist position. Rather, it can be applied to different accounts of knowledge. On the other hand, this new model of the debate is not unfairly unfavourable for absolutism. Instead, it also boosts motivations for advocating absolutism by shifting the focus of the debate from internal absolutism to external absolutism. Precisely put, firstly, this remodelled debate preserves all motivations for being an absolutist that one can hold in the original model of the debate (*e.g.*, linguistic evidence). Besides, it further absorbs some motivations for absolutism that are compatible with internal gradualism, such as epistemicism, moderate K-THRESHOLD, and the credence-one/threshold view of belief. Since Hetherington presents the gradualism/absolutism debate as one between internal gradualism and internal absolutism, these

⁶⁴ In addition, Jaster (2013) explicitly claims that knowledge is a *pure threshold predicate*. For a pure threshold predicate, in each context 'the obtaining standard fixes the conditions that have to be satisfied for a subject to exceed the threshold. Once the threshold is exceeded, the predicate applies.' (Jaster 2013: 321)

motivations are not only unavailable but also inconsistent with the absolutist side within the current model.

The last point that I shall mention without expanding is: this new model can also yield a more fruitful debate that can shed light on many contemporary epistemological issues. For example, the threshold view and the spectrum view of knowledge could lead us to two distinct directions of epistemological investigation. By endorsing K-THRESHOLD, epistemology should strive for locating the threshold for knowledge, or at least, narrowing down the span within which T_k is located. Call this paradigm of epistemological investigation *the threshold paradigm*. In contrast, if K-SPECTRUM is accepted, then we should not even attempt to narrow down the span, as there is no threshold for knowledge at all. Rather, it is more meaningful to figure out indicators of paradigmatic cases of knowledge and not-knowledge so that we can depict the spectrum of knowledge in a more fine-grained manner. Call this paradigm of epistemological investigation *the spectrum paradigm*. Under the threshold paradigm, many advocates⁶⁵ of K-THRESHOLD see it as a primary task for epistemology to give a reductive analysis of knowledge, or, necessary and sufficient conditions for knowledge. That is because a threshold for knowledge will usually serve as an analysis of knowledge, or indicate the necessary and sufficient conditions for an epistemic standing's being qualified as knowledge. On the contrary, under the spectrum paradigm, the anti-reductive-analysis line (see, for example, Zagzebski, 1994; Williamson, 2000) is what we should take. It is in principle impossible to define a spectrum concept 'red' in terms of necessary and sufficient conditions whereby we can distinguish red from any other colour. What is important for comprehending a spectrum concept is not *analysis*, but *characterisation*—characterising what a typical instance of that concept is like and what a less typical instance would be like, and so forth. It is not totally crazy to grasp a

⁶⁵ Admittedly, one can deny that knowledge is reductively analysable but still endorse K-THRESHOLD. The point here is just that a typical characteristic of the threshold paradigm is the pursuit of reductive analysis of knowledge. At least, the threshold paradigm has more motivations to pursue such an analysis compared with the spectrum paradigm. For the spectrum paradigm, a reductive analysis of knowledge should not have been expected at all.

concept without resorting to its analysis or necessary and sufficient conditions. Consider family resemblance concepts—by definition, they have no commonly shared core/essence encompassing their definitions in terms of necessary and sufficient conditions. There can be no reductive analysis for a family resemblance concept, such as ‘games’. Moreover, it is impossible to draw a sharp boundary dividing all members of ‘games’ from all members of ‘not-games’. Admittedly, it is worthy of debate whether knowledge is a family resemblance concept⁶⁶—what I try to illustrate here is just that a spectrum paradigm of conceptual investigation is not only conceivable but maybe also more promising than it appears to be. With these being said, it should be evident that, being based on K-THRESHOLD and K-SPECTRUM, the debate between absolutism and gradualism is not merely a debate around a property (*i.e.*, gradability) of knowledge. More importantly, it also presents us with a competition between two modes of epistemological investigations which could influence the direction that relevant researches should head to.

The fruitfulness of the remodelled debate also embodies in how gradualism and absolutism can be respectively applied to solve epistemological issues. The better one side of the debate can help to solve contemporary epistemological puzzles, the more likely that it would be seen as a better reading of knowledge and thus wins over more support. For example, which view can better answer the value problems of knowledge? Can the Gettier problem gain better resolutions within a gradualist framework? How should we rethink the relation between knowing and understanding?

Important and interesting as they are, this chapter is unable to address these problems due to the limited length. Nevertheless, at this stage, I hope that I have accomplished the primary goal of this chapter—to reconstruct a better framework for the gradualism/absolutism debate. In the remainder of this thesis, I will manage to exhibit the theoretical advantages that gradualism

⁶⁶ Kusch (2011) and Kusch & McKenna (2018) argue that knowledge is a family resemblance concept.

enjoys over absolutism. The last section of this chapter will take the first step—it will be proved that gradualism can offer us a unified account of different types of knowledge, which can explain away the putative asymmetry between knowledge-that and other types of knowledge in terms of the gradability.

4. Towards A Unified Account of Knowledge

The gradualism/absolutism debate is an epistemologically valuable issue that is worthy of further discussions. However, the extant model of the debate interprets the core divergence of gradualism and absolutism in a problematic way, and thus succumbs to the two-sense problem. I suggest that we should remodel the debate by focusing on the competition between the threshold reading and the spectrum reading of knowledge, given that it can help us to construct a more philosophically meaningful and fruitful debate. I am far from believing that the considerations developed in this chapter make a conclusive case for the debate between gradualism and absolutism. Instead, I hope that this chapter could be seen as a restart rather than a conclusion of the relevant debate.

There are many types of knowledge in terms of the object to be known. The focus of this thesis is knowledge of a proposition, *i.e.*, knowledge-that. Besides, knowledge-who, knowledge-where, knowledge-why and the like constitute knowledge-wh. Knowledge of how to do something is called knowledge-how⁶⁷. Similar taxonomy also comes in the form of propositional knowledge, objectual knowledge, practical knowledge, and so on. Reductionism argues that different types of knowledge can be reduced to one single basic type of knowledge. For instance, intellectualism holds that knowledge-how can be reduced to knowledge-that, while the radical anti-intellectualism argues that knowledge-that is a species of knowledge-how.

⁶⁷ Similarly, Hetherington invented a jargon 'how-knowledge' which captures not only the intension of knowledge-how, but also a broader (while somehow vaguer) sense of 'how it is that *p*'.

However, among those varieties of knowledge, knowledge-*wh* and knowledge-*how* are ordinarily taken to be gradable concepts (as we have discussed in the first chapter). In contrast, the traditional absolutist view claims that proposition knowledge is significantly heterogeneous—it is non-gradable. This becomes an apparent challenge to reductionism. However, the absoluteness of knowledge-*that* suggests that there is something fundamentally different between ungradable knowledge-*that* and gradable knowledge-*how*, which makes the reductionism dubitable. This also constitutes a classic anti-intellectualism approach in the literature (see Ryle, 1949; Sgaravatti and Zardini, 2008; Michaelis, 2011; Bengson and Moffett, 2011; Wiggins, 2012; *etc.*). Radical-intellectualism also faces the same difficulty. Hetherington (2005) evades this difficulty by combining his radical-intellectualism with epistemic gradualism, to the effect that both knowledge-*that* and *how*-knowledge are gradable⁶⁸.

Let me clarify the dialectical state of play now. Although I do not think gradualism has to imply reductionism, it is still necessary for anyone who advocates reductionism to face up to the gradualism/absolutism debate. That is because if knowledge-*that*, as opposed to other varieties of knowledge, is not gradable, then this asymmetry becomes an important explanandum that reductionism has to explain properly. Note that this challenge is not only for reductionism. For those who endorse both absolutism and anti-reductionism, it still remains the problem that what makes knowledge-*that* so different. How should we explain the asymmetry between knowledge-*that* and other types of knowledge in terms of their gradability? Call this problem *the asymmetry problem*. This problem can be further divided into more sub-problems, for example: What is the heterogeneous trait that propositional knowledge has, while other types of knowledge do not? When taking the genealogy of knowledge into consideration, what is the essential difference between the function of knowledge-*that* and that of other types of knowledge? Does this

⁶⁸ However, as we have seen in the previous chapters, his attempt at basing gradualism on *how*-knowledge is not successful. In particular, the reductionist assumption should not be seen as necessary for defending gradualism.

difference really exist? Is there another distinct genealogy storyline about the formation of our concept of know-how/know-wh, compared with that of know-that? To answer these questions, absolutists are prone to complicate their explications so that the alleged asymmetry can be plausibly addressed. For example, they need to distinguish the nature of knowledge-that from that of knowledge-how/knowledge-wh, and provide extra explanations about how this difference came into being. Even if these tasks can be accomplished, this will perforce sacrifice the simplicity of explication. Therefore, an absolutist presupposition would encumber both reductionists and non-reductionists (especially reductionists).

The simplest while plausible solution to this problem is to abandon the absolutist position—not only the beyond-threshold absolutism, but also the threshold absolutism (Hetherington's solution failed partially because he just abandoned the former). By embracing the gradualist account of knowledge, we find the alleged asymmetry disappears as different varieties of knowledge are all gradable. There is no more 'heterogeneous trait' that needs to be explained. This leads us to a unified theory of knowledge where the gradability of various types of knowledge are homogeneous and symmetrical. In contrast with absolutism that brings on heterogeneity and asymmetry, a gradualist account of knowledge is more concise, coherent and simple.

One may question whether a unified account of knowledge is worthy of pursuing. Maybe this pursuit is mistaken from the outset, if knowledge is a family resemblance concept as Kusch (2011) and Kusch & McKenna (2018) suggest. If that is the case, perhaps we shall stop craving for generality or a commonly shared essence for different members of the knowledge family. My response is that, gradualism-based unified account of knowledge is compatible with the idea that knowledge is a family resemblance concept. That is because, the family resemblance view only rejects the existence of a single *core/essence* commonly shared by each type of knowledge. It does not imply that members of knowledge family cannot have anything in

common. For example, for the family of numbers, all types of number have at least one thing in common: they are all numbers. Family resemblance does not refuse commonly shared characteristic, but just commonly shared core/essence that can encompass the definition of knowledge in terms of necessary and sufficient conditions.

In his *Blue Book*, speaking of family resemblance and philosophers' inclination of 'carving for generality', Wittgenstein writes:

'We are inclined to think that there must be something in common to all games, say, *and that* this common property is the justification for applying the term "game" to the various games' (Wittgenstein, *Blue Book*, 1958: 17; emphasis mine)

It is important that we do not dismiss the latter half of the sentence after 'and that'. For a general term that counts for a family resemblance concept, what Wittgenstein denies is not that there is a common property *simpliciter* shared by all instances, but that there is a common property that justifies why these instances are all instances *of* the very general term⁶⁹. If the thesis of family resemblance just indiscriminately denies any common property, then the thesis is clearly false. Obviously, there are some traits (albeit trivial) common to all types of games, say, they are all human activities, they are all 'playable' for their game-players, *etc.* A sympathetic reading of the thesis of family resemblance should better consider what question this thesis aims to answer. In *Philosophical Investigations*, Wittgenstein starts the talking of family resemblance by responding to an imagined critique:

'You take the easy way out! You talk about all sorts of language games, but have nowhere said what the essence of a language-game, and hence of language, is: what is common to all these activities, and *what makes them into language or parts of language.*' (PI, § 65; emphasis mine)

⁶⁹ I am not the only one who holds this reading. See also M.A. Simon (1969), Jert (1995), *etc.*

It should be evident that the issue that Wittgenstein intends to address is: what makes us identify something as belonging to a kind or a general term. What makes us call all instances of language as 'language'? Thereupon, his answer is: there is no such a common essence *that can make us apply the same term to all its instances*:

'Instead of producing something common to all that we call language, I am saying that these phenomena have no one thing in common *which makes us use the same word for all*, — but that they are related to one another in many different ways. And it is because of this relationship, or these relationships, that we call them all "language".' (PI, § 65; emphasis mine)

Taking the context of PI into consideration, we have no reason to dismiss Wittgenstein's restriction on the non-existent common property—it should be a property that serves as an essence, to wit, a sufficient and necessary condition for membership in the extensions of a general term.

With these being said, it is not hard to see why the commonly shared gradability among different members of 'knowledge' does not contradict the idea of family resemblance. Gradability does not serve as an essence for knowledge that can sufficiently and necessarily determine extensions of knowledge. It is just a commonly shared characteristic of different varieties of knowledge. On the contrary, the existence of such a common property can help us better understand why knowledge-that, knowledge-how and knowledge-wh belong to the same *family*.

So now we have seen that epistemic gradualism can provide us with a simple unified theory of knowledge in terms of the gradability. Is there any competing proposal that argues the opposite? Recently, Carlotta Pavese (2017) has proposed a competing account arguing that knowledge-how is also *ungradable* in its nature, despite our ordinary talk about gradable know-how. According to Pavese, her project aims 'to argue for the compatibility of the phenomenon of gradability with the absoluteness of the state ascribed by

know-how ascriptions' (2017: 350). Pavese suggests that linguistic phenomena of ordinary epistemological talks are insufficient to reveal the conceptual nature of an epistemic state. I agree with Pavese on this point (see Chapter Two). In my case, knowledge-that is gradable, despite ordinary talks about ungradable know-that; in Pavese's case, knowledge-how is ungradable, despite ordinary talks about gradable know-how. Secondly, Pavese and I both strive for a unified account of knowledge that can uniformly explain knowledge-how and knowledge-that. The motivation of Pavese's project is to defend the *intellectualism* of knowledge-how, a view that knowledge-how is reducible to knowledge-that, from the apparent asymmetry of their (knowledge-how and knowledge-that) gradability. If Pavese's project succeeds, it would also bring us to a unified solution to the asymmetry problem—both knowledge-how and knowledge-that are ungradable.

Whereas, I do not think Pavese's competing project can prevail over mine. First, her ungradable reading of knowledge-how is based on the intellectualism of knowledge-how. Pavese advocates that one's knowledge of *how to ϕ* can be understood as one's knowledge of a practical answer to the question 'how to ϕ '—and the practical answer can be propositional. Hence one's knowledge-how is equivalent to one's propositional knowledge of a (bunch of) proposition(s) expressing the practical answer to the question 'how to ϕ '. Given that, Pavese proposes analyses as follows: Quantitative talk of gradable know-how expresses just a matter of knowing different extents of the practical answer. Qualitative talk of gradable know-how such as 'A knows how to ϕ better than B does' just means that A's practical answer to 'how to ϕ ' is better than B's (see Pavese, 2017). All in all, what is gradable is just the practical answer, rather than one's (propositional) knowing that answer. One's propositional knowledge of the answer, and thus one's knowledge-how, is still an absolute matter. Clearly, Pavese's analyses are premised on the intellectualism, which has invited a historic debate where numerous philosophers have developed various forms of anti-intellectualism

(see Ryle 1949; Roland 1958; Carr 1981; Noë 2005; Hetherington 2006, 2011; Cath 2011; *etc*). I do not plan to expand the debate here, but just mean to call for readers' notice that it seems that in regard to whether know-how is exhaustively reducible to know-that, the foundation of Pavese's project is dubitable. For example, Carter & Pritchard (2015) argue that knowledge-how is not a kind of knowledge-that because the former does not share an important epistemic property with the latter. The idea here is that knowledge-how is compatible with a variety of Gettier-style epistemic luck (in particular, environmental luck), while knowledge-that is not. Furthermore, this verdict is supported by a recent experimental study conducted by Carter *et al.* (2019), which reports that people are even more inclined to attribute knowledge-how in cases where epistemic luck is present; in contrast, knowledge-that attributions are not sensitive to epistemic luck.

Moreover, even following the intellectualist line, Pavese's conclusion is still drawn too quickly. That is because Pavese's analyses can at most show that there is a way to explain linguistic evidence of gradable know-how talk without imperilling the propositional essence of knowledge-how. That is, we talk of know-how as if it is gradable, but in its nature, know-how is still a species of know-that. However, Pavese takes it for granted that knowledge-that is ungradable. She provides no argument defending her absolutist account of propositional knowledge, which is exactly what my proposal urges to reject. In other words, 'whether know-that (rather than know-how) is gradable' is the central dispute, where Pavese talked less. In addition, it is unclear why the quality (or say, the gradability) of one's practical answer cannot reflect the quality (gradability) of one's knowledge-how. It seems to be self-evident for Pavese that the quality of one's practical answer is fundamentally distinct from the quality of one's corresponding knowledge. But this might not be that self-evident. Why cannot gradualist of knowledge-how argue that one's knowledge-how-to- ϕ is gradable just in terms of the quality of one's practical answer to 'how to ϕ '? Especially, drawing on the *end of inquiry* account, the primary function of knowledge-how is just to signal that

our cognitive inquiry towards ‘how to do something’ can be ended as a good-enough answer is acquired. Hence it makes good (if not perfect) sense that the quality of one’s knowledge-how is subject to the quality of the answer acquired. This can better explain why we need knowledge-how and when the inquiry can end—for absolutists, there is an absolute cut-off point where the inquiry could and should end⁷⁰; while for gradualists, the end-point can be relative to different situations. I find the latter narrative seems to match our daily epistemic practice better⁷¹.

Now we have discussed the competing unified theory and noticed that it lacks competitiveness. Endorsing the gradualist account of knowledge remains to be the most promising unified theory that is of the merit of simplicity and efficiency.

5. Concluding Remarks

In Chapter Three, we analysed the defects of Hetherington’s two versions of gradualism. While in this chapter, we further discussed the defects of the current model of the debate between gradualism and absolutism that is based on Hetherington’s problematic gradualist proposals. In order to avoid those defects, this chapter has explored a new way to reconstruct the model of the debate by proposing the distinction between K-THRESHOLD and K-SPECTRUM. With a better model of the debate in play, we can now further build up motivations for embracing gradualism. My strategy is: if gradualism can help to address significant epistemological issues that can hardly be

⁷⁰ What is worse, Pavese’s absolutist intellectualism of knowledge-how does not seem to provide an explanation about when the inquiry can come to an end. That is because her account of know-how only emphasises on one’s (propositionally) knowing a practical answer, but says nothing on how the quality of that answer influences the truth value of the corresponding knowledge ascription. After all, this influence is what Pavese intends to cut off from the nature of know-how, so that the quality of the answer does not affect the quality of the knowledge.

⁷¹ Consider it: is there actually a cut-off point where one can be recognised as knowing how to play the piano?

solved within an absolutist picture, then gradualism should be taken to be preferable to absolutism.

I have to admit that Hetherington also took a similar strategy to advocate gradualism: he also claimed to show gradualism's theoretical merits in terms of its applications in scepticism, the Gettier problem, a unified account of knowledge, and even the McKinsey paradox (see Hetherington 2001; 2011, ch. 5). So it seems that Hetherington has already shown gradualism's advantage in the same manner that I am about to do. Why is my attempt still needed? The answer is straightforward: the success of Hetherington's comparison between gradualism and absolutism is based on a premise that his gradualist proposals hold water. However, as we have seen in Chapter 3, both versions of his gradualist proposals have significant defects, and therefore we need a new account of gradualism. Defects of Hetherington's gradualist accounts of knowledge will also affect his arguments showing gradualism's theoretical advantages. In Chapter 3, it has been shown that his application of gradualism in the Gettier problem is unpersuasive, and how problems of his how-knowledge-based account of gradualism regrettably hinder his reductionist unified account of knowledge. In Chapter 6, we will further see why his gradualist anti-sceptical proposals are also unsuccessful.

We have seen that the first valuable application of gradualism is that it can offer us a unified account of knowledge that enjoys the merits of simplicity and efficiency. The next chapter will illustrate several more possible applications of gradualism.

Chapter Five: Gradualism and Explanatory Power

Abstract: This chapter purports to show gradualism's theoretical merits by exploring its applications in solving two important epistemological issues. The first section focuses on epistemic luck. It is controversial whether a specific type of epistemic luck, namely, *environmental luck*, is incompatible with knowledge. Goldberg (2015) proposes a new account of epistemic luck that is based on the notions of *epistemic entitlement* and *adequate explanation* (in contrast to Pritchard's safety-based account), and is alleged to be able to account for the two conflicting intuitions regarding environmental luck's compatibility with knowledge. I will argue that Goldberg's account of epistemic luck suffers from two defects: 1) it will end up denying the notion of environmental luck; 2) it cannot accommodate the intuition that environmental luck is compatible with knowledge. Instead, gradualism can reconcile the debate regarding environmental luck in a more satisfactory manner. The second section will explore gradualism's application in accounting for faultless disagreements.

Faultless disagreements on knowledge ascriptions have two seemingly incompatible conditions: 1) two (or more) competent epistemic peers are genuinely disagreeing with each other on the truth value of a knowledge ascription proposition (call it DISAGREEMENT), and 2) none of them is genuinely making a mistake (call it FAULTLESSNESS). How shall we reconcile the two conditions? Call this *the puzzle of faultless disagreement*. Three mainstream solutions to the puzzle (namely, indexical contextualism, non-indexical contextualism, and non-indexical relativism) will be examined. The three extant proposals all succumb to a dilemma: they would either sacrifice FAULTLESSNESS or DISAGREEMENT. In contrast, gradualism can better

accommodate both conditions by interpreting knowledge as a spectrum concept analogous to 'red', 'warm', or 'bright'.

The previous chapter has introduced a new model of the debate for gradualism and absolutism. We have seen that a gradualist theory of knowledge based on K-SPECTRUM could lead this debate, as well as gradualism itself, to a more fruitful prospect. Gradualism's advantages over absolutism can be reflected by its theoretical fruits. K-SPECTRUM's application in resolving the putative asymmetry problem is an instance. This chapter will further explore gradualism's implications in contemporary epistemology. It will be argued that gradualism enjoys higher explanatory power than absolutism in solving two knotty issues in the contemporary epistemological debates that involve epistemic luck and faultless disagreements.

1. Gradualism and Epistemic Luck

1.1: Two intuitions

Gettier-style counterexamples famously challenged the JTB account of knowledge. Any post-Gettier analysis of knowledge is expected to be able to avoid the Gettier-style challenge and explain away Gettier-style cases. The influential anti-luck epistemology attributes protagonists' lack of knowledge in Gettier-style counterexamples to knowledge-undermining epistemic luck. It is argued that knowledge cannot be the result of mere epistemic luck, while in Gettier-style cases, the protagonists' beliefs are just luckily true.

Pritchard (2009a, 2009b, 2012a) distinguishes two different types of epistemic luck. The first type is *intervening luck*. A knowledge-apt belief of a fact is expected to be true *because* it captures the target fact. While in cases such as Gettier's original coins-in-pocket case and Ford-or-Barcelona case (see Gettier 1963), and Chisholm's sheep-in-the-field case (see Chisholm

1966), there is a kind of *intervening luck* that intervenes between the subject's belief and the target fact. In those cases, the protagonists' beliefs turn out to be true not in virtue of their capturing the target fact by manifesting their cognitive abilities, but merely in virtue of the intervening luck.

In contrast, *environmental luck* refers to the type of luck that is involved in fake-barn-county style Gettier cases (see Goldman 1976) where there is nothing intervening between the belief and the target fact. For example, in the fake barn case, the protagonist (name her 'Barney') does see a real barn and her belief that 'there is a barn in front of me' does capture the fact that she sees a real barn. However, her belief is still luckily true as she is in a bad epistemic environment where she could have easily formed a false belief by mistakenly taking a barn façade for a real barn. A similar case of environmental luck is Pritchard's fire-officer case (see Pritchard 2010), where the protagonist occasionally encounters a real fire officer among a group of cleverly disguised fake fire officers. The protagonist forms a true belief regarding the cause of a fire emergency by virtue of the real fire officer's testimony. However, this belief is just luckily true because our protagonist could have easily talked with a fake fire officer who would only lie. The absence of knowledge in those cases is not due to the absence of cognitive achievement, but due to the inhospitable epistemic environment⁷².

The safety-based account of epistemic luck accounts for these two alleged knowledge-undermining veritic epistemic luck in terms of the notion of *safety*. That is, a true belief is safe *iff* it could not have easily been formed on the same basis and yet been false (see, for example, Sosa 1999; Williamson 2000; Pritchard 2007). True beliefs in Gettier cases fail to constitute knowledge because they all could have easily been false, no matter the unsafety is caused by intervening factors or environmental factors.

⁷² Interestingly, Gendler and Hawthorne (2005) report that many people find there is a tendency to ascribe knowledge in a highly similar case that they dub WATCH OUT: a subject forms a true belief regarding what time it is by asking a reliable truth-teller who happens to be surrounded by a roomful of compulsive liars.

It is widely accepted that knowledge cannot be the result of mere epistemic luck. Also, it is relatively uncontroversial that classic Gettier-style luck is knowledge-undermining. In contrast, there has been a debate on whether environmental luck is incompatible with knowledge as Pritchard claims. Many, albeit not mainstream, epistemologists share the intuition that environmental luck *is compatible* with knowledge, thus Barney *does know* that there is a real barn in front of her. Prominent advocates of this intuition include Hetherington (1999, 2001), Gendler and Hawthorne (2005), Lycan (2006), Sosa (2007, 2009), Turri (2011b), and Colao, Buckwalter & Stich (2012).

By way of illustration, Hetherington (2001) argues that in the fake barn case, Barney (poorly) knows that there is a barn in a very bad way, however, her true belief still counts as knowledge, given that mere true belief is sufficient for constituting the minimal knowledge. Sosa (2007) holds that Barney's belief in the fake barn case is apt, to wit, her belief is true because of her reliable cognitive competence (e.g., perceptual competence). Thus, her apt belief suffices to count as a piece of *animal knowledge* despite the luckiness. We find the quality of Barney's knowledge not very satisfactory because it falls short of *reflective knowledge*, which is a higher-level of knowledge that 'requires also an apt apprehension that the object-level perceptual belief is apt' (Sosa, 2007: 108). Drawing on Sosa's AAA-model (accurate, adroit, apt) of knowledge, Turri (2011b) adds one more condition for knowledge—to constitute knowledge, a belief should not only be *apt* (true *because of* competence), but also *adept* (a true belief that *manifests* competence)⁷³. In accordance with this analysis of knowledge, Barney's true belief does manifest her relevant cognitive competence, and thus should be recognised as knowledge. As for the prevailing intuition that knowledge is absent in the fake barn case, Turri directly denies that he shares this intuition. To persuade the readers, Turri invites us to consider a modified case: Bad Henry

⁷³ For a detailed discussion of the aptness/adeptness distinction, see Turri (2011b).

maliciously destroys a barn (that he firmly believes to be a barn) by a bazooka, without being aware that he is in a fake barn county where the only real barn that he destroyed is surrounded by indistinguishable barn façades. Turri holds that it should be more intuitive that Bad Henry did know that he was destroying a real barn as he pulled the trigger. However, just like Turri lacks the intuition that Barney *lacks* knowledge, advocates of environmental luck's incompatibility of knowledge might also lack Turri's intuition that Bad Henry *possesses* knowledge. Incompatibilists can simply insist, without being stubbornly steadfast, that Bad Henry's belief is just luckily and unsafely true and hence falls short of knowledge. It is irrelevant whether he destroyed the barn or just saw the barn.

In summary, we seem to reach an impasse in the debate. What worsens this impasse is, from an experimental-philosophy perspective, Colao, Buckwalter & Stich (2012) discover that the Turri-style intuition is held by the majority of respondents. So, when the two uncompromising conflicting intuitions pose an apparent impasse, is there any possibility of putting forward a reconciliation that breaks the deadlock? Moreover, how shall we account for the existing divergence? That is, each side of the dispute seems to hold a reasonable intuition—notwithstanding, the two reasonable intuitions turn out to constitute a contradiction. Is there any plausible explanation for this paradoxical phenomenon?

1.2: Goldberg's Reconciliation

We have seen that there seems to be a stalemate in the debate about the compatibility between environmental luck and knowledge. The two sides of the debate hold two conflicting intuitions and neither side can convince the other to abandon their intuition. A reconciliation would be valuable if it can break the current deadlock of dialogues and account for the seemingly paradoxical phenomenon that two reasonable intuitions reach two conflicting conclusions regarding Barney's epistemic status. Few researches in the

literature attempted to reconcile these two apparently irreconcilable intuitions. An exception is Goldberg (2015).

Goldberg argues that Pritchard's safety-based account of epistemic luck cannot fully capture the sense that knowledge cannot be lucky, since there can be lucky true beliefs *that are safe*. He cites some influential counterexamples to the safety account of epistemic luck, for instance, Lackey's SOUTHERNMOST BARN case:

'SOUTHERNMOST BARN [W]hile entering a Midwestern farming community on her cross country drive, Janice looked at the first barn that she saw, which was on the southernmost end of the field, and formed the corresponding belief 'There is a barn'. As it happens, the barn she saw is the only real one, surrounded by barn facades that members of this community have placed in the field in order to make their town appear prosperous. However, as a matter of strict and unwavering policy, the members of this community always place their only real barn on the southernmost end of their land, since this is where traffic first enters their town. Moreover, thirty years earlier, Janice had lived in a house on the southernmost end of this field in the precise location of the one real barn. Because of her deep interest in her childhood roots combined with the brief period during which she can safely take her eyes off of her driving, she would invariably have looked at only the particular place in the field where the real barn exists.' (Lackey 2006: 288)

Goldberg argues that, in this modified fake-barn-county case, Janice's belief should be seen as safe if we adopt Pritchard's modal account of safety.

Pritchard defines safety as follows:

'S's belief is safe *iff* in most nearby possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world the belief continues to be true.' (Pritchard 2007: 281)

For the safety-based account of epistemic luck, a true belief is lucky *iff* it is not safe. However, in the SOUTHERNMOST BARN case, in most nearby possible worlds Janice would always invariably have looked at the southernmost end of the field where the real barn stands, thus her belief that ‘there is a barn’ would continue to be true in those worlds. According to the modal account of safety, her belief is safe. Nevertheless, Lackey, as well as Goldberg, argues that Janice’s true belief still suffers from epistemic luck. After all, it is purely a matter of luck that, first, residents there decided to put the only barn in the precise location where used to be the very house that Janice once lived in; second, Janice was so interested in her childhood roots that she would invariably look at the real barn rather than those barn façades aside. So there is a specific kind of epistemic luck, which is not associated with the unsafety of Janice’s belief, but is rooted in the lack of explanatory connection between her belief’s being safe and her epistemic agency. That is, the safety of Janice’s true belief has nothing (or at least little) to do with the manifestation of her relevant cognitive competence.

Given these safe-but-lucky counterexamples⁷⁴, Goldberg proposes a new account of epistemic luck that is based on the notion of *epistemic entitlement*. Goldberg defines epistemic entitlement in this way:

‘[A] subject S is epistemically entitled to the proposition that *p* just in case it would not be epistemically improper for S to presuppose that *p* in the course of belief-formation or deliberation.’ (Goldberg 2015: 281)

According to Goldberg, here are two varieties of epistemic entitlement: earned and unearned. One’s entitlement to a proposition *p* is *earned* just in

⁷⁴ Goldberg’s motivation for rejecting the safety-based account of epistemic luck is more sufficient than solely a SOUTHERNMOST BARN case. SOUTHERNMOST BARN is just a counterpart of Goldman’s fake-barn-county style luck. In fact, he argues that there are also counterexamples mirroring Gettier’s original Ford-or-Barcelona style luck and Chisholm’s sheep-in-the-field style luck (for the former, see Hiller & Neta 2007: 308; for the latter, see Goldberg 2015: 277-278). The reason why I only select to present SOUTHERNMOST BARN is that it is an example involving environmental luck, which matches the purpose of this chapter most.

case *p* is either doxastically or propositionally justified for him/her. There are also some unearned entitlements relevant to one's proper reliance on *routes to knowledge*, *i.e.*, methods (such as perception, memory, and testimony) that when employed under normal conditions and working properly, often yield knowledge. Roughly speaking, a route to knowledge is required to be a reliable belief-forming method that often results in the acquisition of knowledge. Goldberg claims that when one properly relies on an epistemic method that qualifies as a route to knowledge, one is presumptively (though defeasibly) entitled 'to presuppose such things as that the method is reliable when functioning properly *under normal conditions*' (Goldberg 2015: 281; emphasis mine).

After introducing these central terminologies, Goldberg formulates his account of epistemic luck as follows:

'EL S's true belief that *p* (acquired through method *M* in circumstances *C*) suffers from knowledge-undermining epistemic luck *iff* there is no available explanation for the fact in question—how it came to pass that S's employment of *M* in *C* resulted in a belief that is true—in terms of those propositions to which S has an earned entitlement, together with those (KR-) propositions to which S is presumptively epistemically entitled in virtue of her reliance on *M* in *C*.' (Goldberg 2015: 285)

Briefly put,

'EL_{Brief} S's true belief that *p* suffers from epistemic luck when the *veritic explanandum* is not adequately explained in terms of S's corresponding entitlement propositions.' (Goldberg 2015: 285)

Two more notions to unpack. First, 'entitlement propositions' refer to propositions that one is entitled. These include those unearned entitlements to propositions such as 'the environmental conditions are normal'—Goldberg calls this *the normal condition proposition*; and 'my cognitive faculties are

functioning properly’—Goldberg calls this *the proper cognitive functioning proposition*; and ‘my perceptual faculties are reliable when functioning properly in normal conditions’—Goldberg calls this *the reliability proposition*. Those propositions are entitled when one is properly relying on a route to knowledge. They are altogether named *KR-propositions* by Goldberg. Second, as for what counts as an *adequate explanation*, Goldberg admits that he does not have a clear criterion to offer—whether an explanation adequately explains its explanandum tends to be a matter of intuitive verdicts.

So how can this new account of luck be applied to solve the Gettier problem? By way of illustration, in the Ford-or-Barcelona case, the protagonist Smith forms a true belief that ‘Jones owns a Ford, or Brown is in Barcelona’. He believes in the first disjunct of this proposition because he overhears a conversation discussing Jones’ Ford car. By the rule of disjunction introduction, he concludes that the whole disjunction must be true, although he does not have any information about Brown’s location. Unbeknownst to Smith, the first disjunct of the proposition turns out to be false, while the second disjunct is surprisingly true. Hence, Smith’s true belief is a result of epistemic luck. Goldberg can account for this luck by diagnosing that, the *veritic explanandum* of this case, *viz*, how it came pass to that Smith’s overhearing and inferring resulted in his belief’s being true, cannot be adequately explained by propositions that he was entitled to. After all, the actual reason why Smith’s belief turned out to be true was that he constructed a disjunction by randomly introducing an accidentally true disjunct. But Smith was not entitled to presume that his randomly introducing a disjunct could yield knowledge—it is not a reliable route to knowledge. He was at most entitled to presume that if the interlocutors in a conversation that he overheard were reliable, then he could take their testimonies at face value—but this obviously does not suffice to explain how it came pass to that his belief was true. As a result, according to the entitlement-based account of epistemic luck, Smith’s belief is only luckily true.

It seems that Goldberg's account works smoothly when accommodating *intervening luck* that occurs in cases like Ford-or-Barcelona. But what about *environmental luck*? Goldberg alleges that one virtue of his account is that it can address both intuitions regarding whether Barney knows in the fake barn county. For the knowledge-absence intuition adjudicating that Barney lacks knowledge, Goldberg's entitlement-based account of luck can explain that, this is because Barney's epistemic success is due to the fact that she looks at the only real barn rather than barn façades around by coincidence. Her entitlement propositions cannot explain why she happens to glimpse at the only real barn. Hence Barney's true belief is lucky and thus falls short of knowledge.

As for the knowledge-presence intuition adjudicating that Barney does know that there is a barn, Goldberg's explanation is that it can be interpreted that Barney's epistemic success can be fully explained in terms of her manifesting her perceptual competence, which is a reliable route to knowledge. Barney is entitled to presume that her perceptual faculty is functioning well (which is true), and that her properly functioning perceptual faculty is reliable (which is true as well). In addition, her belief that 'there is a barn' is actually justified by her seeing that barn, thus she also enjoys the relevant earned entitlement. Putting these all together, her true belief can be adequately explained. Goldberg admits that there is something abnormal regarding Barney's epistemic environment—her environmental conditions are abnormally inhospitable. However, he suggests that advocates of the knowledge-presence intuition can argue that 'this abnormality is irrelevant to how it came to pass that Barbara's reliance on perception in this circumstance resulted in a belief that is true' (Goldberg 2015: 289). After all, Barney is not looking at a fake barn, and what explains her epistemic success is but her seeing and recognising the real barn, together with her corresponding unearned entitlements in that circumstance. The reason why she happens to see the sole real barn rather than other fake barns is irrelevant to the explanation of

her epistemic success. We can better comprehend this irrelevance by considering a detective who knows a key clue by accidentally glimpsing at a corner of the room where the murder was conducted. Drawing on this key clue, the detective thereby reasons out the identity of the murderer. The detective knows the clue, and he knows who the murderer is. According to Goldberg's suggested diagnosis, for people who attributes knowledge to the detective, the explanation of the detective's epistemic success is irrelevant to why he happens to cast his eyes at *that* very corner rather than other corners or any place else. His epistemic success can be adequately explained by his perceptual contact with the clue and the consequent recognizing-inferencing process. Similarly, in the fake-barn-county case, we do not need to explain why Barney happens to see the real barn in order to adequately explain her epistemic success. The upshot is, by taking this line of explanation, Barney has a true belief that is not luckily true, and thus one can reasonably attribute knowledge to her. In other words, in the fake barn case, there is no epistemic luck, and hence there is knowledge.

I doubt whether Goldberg's account of epistemic luck can properly explain away the knowledge-presence intuition. As we have seen, he accommodates this intuition by arguing that Barney's true belief, *contra* what Pritchard and most anti-luck epistemologists would predict, does not suffer from epistemic luck. However, this explanation faces the risk of eliminating the concept of *environmental luck* thoroughly. That is not only because the fake-barn-county style cases are seen as paradigmatic instantiations of environmental luck in the present literature, but also because that the notion *environmental luck* is exactly used to refer to epistemic luckiness caused by abnormal environmental factors. If Goldberg's explanation for the knowledge-presence intuition holds water, then one could easily exclude all abnormal environmental factors from the full explanation of one's epistemic success—and hence the judgment of whether epistemic luck involves. Abnormal environmental factors are irrelevant to the identification of epistemic luck. The corollary is that, even in cases where the environmental conditions are

abnormal (such as Barney's case), there is no such thing as environmental luck. As a result, there would be no cases for environmental luck—it is an empty and inconceivable concept.

Furthermore, even if we endorse Goldberg's account of epistemic luck and bite the bullet that there is no such thing as environmental luck⁷⁵, the difficulty still remains. That is because, if we examine the entitlement-based account of luck strictly, it would be found that this account is just incompatible with the knowledge-presence intuition. It would fail to explain why Barney has knowledge in the fake barn county, as Barney's relevant *entitlement propositions* cannot explain why her belief is true. Remember how Goldberg characterises KR-propositions:

‘Suppose that conditions are, and seem to S to be, *normal*..... *In these circumstances*⁷⁶, I submit, S ... enjoys an unearned entitlement to the proposition that environmental conditions (which seem to her to be normal) are normal—call this the *normal condition* proposition. She enjoys an unearned entitlement to the proposition that her cognitive faculties (which seem to her to be functioning properly) are functioning properly—call this the *proper cognitive functioning* proposition. She enjoys an unearned entitlement to the proposition that her perceptual faculties are reliable when functioning properly in normal conditions—call this the *reliability* proposition.’ (Goldman 2015: 282-283)

It is obvious from this quotation that a subject's enjoying entitlement to various KR-propositions is premised on her being in a *normal condition*, which requires a normal epistemological environment. This is understandable, because in an unfriendly epistemological environment, for example, an environment where the lighting is bad, the air is hazy, or there exists a lot of misleading distractions that could have easily lead the subject

⁷⁵ In that case, the alleged virtue of Goldberg's theory will not be that it is able to reconcile conflicting intuitions regarding whether *environmental luck* is incompatible with knowledge, but only that it can reconcile conflicting intuitions regarding whether knowledge occurs in the fake-barn-county-style cases.

⁷⁶ Emphasis mine.

to a false belief, it would be unreasonable to describe that subject as properly relying on a route to knowledge. Only when the conditions are normal, and the subject's cognitive faculties are functioning well, the presumptive entitlement can be made that the subject is employing a method M that will reliably yield knowledge in the circumstance C. However, it is obvious that the epistemological environment of Barney does not fit the bill—her epistemological conditions are abnormal. That is to say, since the *normal condition proposition* is not entitled to Barney, she is not properly relying on a route to knowledge. Consequently, it will need further explanations on how it comes to pass that one forms a true belief by her improper reliance on a method M (in an abnormal circumstance like fake barn county). Barney's epistemic success will fail to be adequately explained in terms of KR-propositions if we apply the entitlement-based account of luck strictly to the fake-barn-county case.

One might argue that the abnormal epistemological conditions will only deprive Barney of those *unearned entitlements* regarding her proper reliance on a route to knowledge, but she can still enjoy the *earned entitlement* to the proposition that there is a barn as her relevant belief is justified by her seeing that barn. But this cannot help to accommodate the knowledge-presence intuition, as merely the earned entitlement is not sufficient for adequately explaining one's epistemic success. Otherwise, the entitlement-based account of luck would do no help to solve the Gettier problem. That is because a basic trait of Gettier counterexamples is that they are cases where the subject has a *justified* true belief but still fails to know. If the mere justification for one's belief can suffice to constitute an adequate explanation for one's epistemic success, then there would be totally no luck in Gettier cases. This would render the entitlement-based account of luck useless.

Now we have seen that Goldberg's entitlement-based account of epistemic luck cannot satisfactorily help anti-luck epistemology to reconcile the conflicting intuitions regarding whether environmental luck is incompatible

with knowledge. It is noteworthy that Goldberg's strategy only explores two possibilities: 1) there is epistemic luck in the fake-barn-county case, and thus there is no knowledge; 2) there is no epistemic luck in the fake-barn-county case, and thus there is knowledge. Both situations seem to presume that veritic epistemic luck must be incompatible with knowledge. The result is, he would either reach a verdict, just like many other anti-luck epistemologists, that environmental luck prevents knowledge from being ascribed; or that there is no such thing as environmental luck. But there should have been a third possibility that we ought to explore: that there is environmental luck in the fake-barn-county case, yet there is also knowledge. To reconcile the two conflicting intuitions regarding Barney's epistemic status, we have to be able to explain the intuition supporting the third possibility.

1.3: Absolutism's Aporia

However, under an epistemic absolutist framework, it seems to be difficult for anti-luck epistemology to accommodate the third possibility. To reconcile intuitions supporting possibility one and possibility three, one has to reconcile the conclusion that environmental luck is incompatible with knowledge as well as the conclusion that environmental luck is compatible with knowledge. This seems to be a mission impossible if there is a threshold for knowledge. Given K-THRESHOLD, Barney's environmentally lucky true belief would either reach the threshold or fail to reach the threshold—it is an absolute yes-or-no affair. Consequently, we have to reject one of the two intuitions, which means the failure of reconciliation. This might also explain why Goldberg just analysed possibility one and two, as if the third possibility is to be taken into consideration, then there seems to be an irreconcilable contradiction: environmental luck is both compatible and incompatible with knowledge. Call this difficulty in accommodating the two conflicting intuitions about the compatibility between environmental luck and knowledge as *the problem of reconciliation*.

There could be two ways for absolutism to solve this aporia. First, a steadfast absolutist solution to the problem of reconciliation can deny that it is worth our effort to reconcile the two intuitions. It can be argued that there is only one correct intuition regarding Barney's epistemic status, and the opposite intuition is just an irrational hallucination that should be eliminated. Second, a more sophisticated absolutist solution can argue that the concept of environmental luck is not fine-grained enough to address the Barney-style cases. We can break down the notion of environmental luck into more sub-notions to the effect that some types of environmental luck are incompatible with knowledge but some are not. It can even be granted that environmental luck comes in different degrees—nevertheless, there is still a threshold that distinguishes knowledge-preventing luck from knowledge-compatible luck.

Both absolutist solutions are defective. Consider the 'FAKE RING COLLECTION' case put forward by Gendler & Hawthorne (2005). Suppose that a lady has one genuine diamond ring and six indistinguishable facsimiles, and she alternates which ring (among the seven rings, genuine and fake) she wears every day. The traditional Goldman-style anti-luck intuition might arrive at a verdict that this lady's friends will *never* know that she is wearing a diamond ring—after all, on any of the other six days, she would have been wearing one of the facsimiles, and her friends would be unable to tell the genuine one from the fakes. The anti-luck intuition supporting this 'no-knowledge verdict' seems to be strong. But now let us modify the case:

FAKE RING COLLECTION II A rich lady has *six* genuine diamond rings and *one* indistinguishable facsimile, and she alternates which ring (among the seven rings, genuine and fake) she wears every day. Unbeknownst to her friends, she only wears the fake ring every Sunday—and she just chooses Sunday as the fake ring-day in a purely random way.

Compared with the original case where the probability of the lady's wearing a genuine diamond ring is only $1/7$, in this modified case the probability rises to $6/7$ —so does the probability of her friends' beliefs being true. As these probabilities rise, I guess most of my readers would agree that the intuition supporting the no-knowledge verdict is weakened. This weakening will be more salient in this further modified case:

FAKE RING COLLECTION III An incredibly rich lady has 364 genuine diamond rings and *one* indistinguishable facsimile, and she alternates which ring (among the 365 rings, genuine and fake) she wears every day. Unbeknownst to her friends, she only wears the fake ring on the first Sunday of April every year—and she just chooses that date as the fake ring-day in a purely random way. Now the question is: one day (any day except for the randomly picked fake ring-day), her friend sees her wearing a diamond ring (as always) and thus forms the belief that 'she is wearing a diamond ring'. Does her friend's belief count as knowledge?

For this further modified case, it would be more farfetched rather than intuitive to still insist that the ring-collector's friends can never know that she is wearing a genuine ring⁷⁷. The steadfast anti-luck epistemology might explain this decline of intuition in terms of the distance of possible worlds where the ring-collector wears her only facsimile. It might be argued that, compared to the actual world⁷⁸ where her friends see her wearing a genuine ring, the only fake ring world is far from being a 'close possible world'. Hence

⁷⁷ If this change is not obvious enough, I invite readers to imagine the case where the ratio of genuine rings to fake rings becomes 36500:1. In that case, it is fair to say that the lady might at most wear the facsimile once in her life.

⁷⁸ For the sake of argument, we can take one genuine-ring world to be the actual world, as if the actual world is the fake ring one then her friends would of course fail to know due to the falsity, rather than the environmental luck, of their beliefs.

the safety account of knowledge can avoid making the counterintuitive verdict that knowledge is absent in the case of FAKE-COLLECTION III.

However, the only reason for us to see the fake ring world as a distant world seems to be its small probability (only 1/365). If that is the case, then the safety account of knowledge would fail to explain why each other genuine-possible world is not distant, as the probability of each world is equal, to wit, 1/365. Moreover, if the fake ring world is distant due to its small probability, then think about the lottery puzzle where one's belief that 'my ticket will lose' only has a much smaller chance (probably one of a million) to be false. When it comes to the lottery puzzle, advocates of safety theory such as Pritchard (2005) would directly refuse to *exclude* this ticket-winning-world from the category of 'nearby possible worlds' on the mere basis of the odds. That is because one alleged theoretical merit of Pritchard's safety theory is that it can solve the lottery puzzle by explaining that the ticket-owner's belief is unsafe given the existence of *nearby* worlds (despite its extremely low probability) where his ticket wins. Since his belief is unsafe, he does not know that his ticket will lose until he sees the result being revealed. The lottery case involves a nearby possible world because of how the ticket is chosen—it has nothing to do with the odds. The ticket is chosen randomly, so is the date when the lady decides to wear the fake ring. One can even impose one more stipulation to the story: the lady picked the date 'the first Sunday of April' by drawing lots, so the date is chosen as randomly as lottery tickets. Therefore, if the extremely unlikely (in a probabilistic sense) ticket-winning world is also a nearby world, then there is no reason that the fake ring world should be seen as a distant.

As a result, the steadfast absolutist solution to the problem reconciliation would be inclined to be farfetched and counterintuitive when it comes to cases like FAKE RING-COLLECTION III. Another case that might embarrass steadfast absolutists is the ALWAYS-WITH-SOMETIMES case (modified from a similar case from Gendler & Hawthorne 2005). Assume that the lady

in FAKE RING-COLLECTION III case is named as Ms. Sometimes. One day, in a mall, Ms. Sometimes, wearing one of her real diamond rings, walks with her friend Ms. Always who also has a same (real) ring and wears it every day. An observer by chance sees Ms. Always' ring and believes it to be a real diamond ring. The question is: does the observer know that Ms. Always wears a real ring? If the steadfast absolutist solution concludes that the observer does not know in the FAKE RING-COLLECTION III case, then the same conclusion should also apply to this case. That is because the two ladies' rings are indistinguishable from each other, and it is highly likely that the observer just casts his eyes to Ms. Sometimes' ring. It is indeed lucky that the observer did not glimpse at Ms. Sometimes' ring—even if he did, he would not have been able to tell anything different. Thus if environmental luck must be incompatible with knowledge, then it would have prevented the observer from gaining knowledge. However, it is even more counterintuitive to refuse to ascribe knowledge in the ALWAYS WITH SOMETIMES case, compared with FAKE RING COLLECTION III. After all, it is untenable to deny that we know that a reliable ring-wearer is wearing a ring just because she has a friend who sometimes wears facsimiles. While if the steadfast absolutism answers that knowledge exists in ALWAYS WITH SOMETIMES, but not in FAKE RING COLLECTION III, then this asymmetry is odd. The steadfast solution thus suffers from a dilemma.

The sophisticated absolutist solution also has difficulty in explaining FAKE-RING-COLLECTION-style cases. If we are committed to the intuition that in the original FAKE RING COLLECTION case, the collector's friends do not have knowledge because of epistemic luck, then why is it much less intuitive that they still lack knowledge in the case of FAKE RING COLLECTION III? The only salient variable for these two cases seems to be the ratio of genuine rings to fake rings. If by our intuition that knowledge occurs in FAKE RING COLLECTION III, then from which point among the scale of genuine/fake rings ratio does knowledge start to emerge? How many real diamond rings should the lady own so that her friends can possess knowledge? How

frequent should the lady wear her fake ring so that her friends fail to know her wearing genuine rings? There does not seem to be a non-arbitrary cut-off point. This difficulty in drawing a threshold is intensified when it comes to a case as follows:

'FAKE BAR Unbeknownst to its patrons, Awful Alvin's Bar serves genuine gin six days per week - and an undetectable surrogate on Sundays. Tom goes out nearly every night; Dick drinks only after his seminar on Tuesdays; Harry is unpredictable but always spends Sundays at home with his family. The three of them gather at Awful Alvin's on Tuesday night, and each of them orders a gin and tonic. Oscar walks in and asks each one what he is drinking. "That's gin," each replies. Does Tom know that he's drinking gin? Does Dick? Does Harry? And does Oscar know that each is imbibing authentically?' (Gendler & Hawthorne 2005: 338)

The four subjects in this case, more or less, all suffer from environmental luck. Nevertheless, it is unclear whether all of them lack knowledge. Moreover, we find it obvious that our confidence in ascribing or dis-ascribing knowledge varies from person to person. For example, maybe for Tom, we will be more inclined to adjudicate that he lacks knowledge, while for Dick and Harry, we will be more willing to recognise them as knowing, because they enjoy the safer (though perhaps not really safe) epistemic environment. Now, how can the sophisticated absolutist solution draw a threshold of knowledge for the four persons? If we change the stipulations slightly, for example, Dick drinks once a week on a randomly-picked day rather than always on Tuesday, can the threshold still apply? I doubt that absolutists can locate such a threshold without being arbitrary. Even if there does exist such a threshold, the difficulty in finding out (or even just narrowing down the range) its location would be an extremely heavy burden of proof for absolutism. Absolutist solutions to the problem of reconciliation are thus severely undermined.

1.4: The Spectrum of Luck

So how can gradualism help to solve the problem of reconciliation? Before constructing the gradualist solution, let us review those intuition data that we are expected to accommodate.

First, it is intuitive that knowledge cannot merely be a result of epistemic luck. It is also relatively uncontroversial that intervening luck is incompatible with knowledge. Hence a gradualist solution ought to capture this anti-luck intuition.

Second, it is also relatively intuitive to conclude that knowledge is absent in the original FAKE RING COLLECTION case, while similar no-knowledge verdicts would be much more controversial when it comes to those modified FAKE RING COLLECTION cases. Also, in cases like FAKE BAR, the epistemic status seems to vary from person to person. Therefore, the gradualist solution is expected to be capable of explaining the changes of intuition, and manage to answer the four knowledge-ascription questions raised in the FAKE BAR case.

Third, as we have seen in the last section, absolutism faces the difficulty in drawing a threshold that distinguishing knowledge-preventing luck and knowledge-compatible luck. While if there is no such a threshold, then how can we capture the anti-luck intuition meanwhile avoiding the absolutist aporia? It would be ideal if the gradualist solution can reach a balance.

With these desiderata being said, now I propose to establish a gradualist account of anti-luck epistemology drawing on Goldberg's entitlement-based version. Call this gradualist account of epistemic luck G-EL:

G-EL (1) S's true belief that p (acquired through method M in circumstances C) suffers from *knowledge-undermining* epistemic luck *iff* there is no *adequate* explanation for the fact

in question—how it came to pass that S’s employment of M in C resulted in a belief that is true—in terms of S’s entitlement propositions (*i.e.*, those propositions to which S has an earned entitlement, and those KR-propositions to which S is presumptively epistemically entitled in virtue of her reliance on M in C).

(2) The less adequate the available explanation is, the more knowledge-underminingly lucky S’s belief is.

(3) S suffers from *knowledge-preventing* epistemic luck *iff* the corresponding explanation is farfetched.

Although G-EL largely inherits Goldberg’s characterisation of epistemic luck, there are still two salient differences. First, G-EL reads the notion ‘knowledge-undermining’ in a way differing from Goldberg’s EL. For Goldberg, ‘knowledge-undermining’ means ‘preventing knowledge from occurring’—if the quality of knowledge is undermined, then there is no knowledge. In contrast, G-EL distinguished ‘knowledge-undermining’ epistemic luck from the ‘knowledge-preventing’ one. The latter refers to epistemic luck that is incompatible with knowledge and would always prevent knowledge from being ascribed; while the former refers to epistemic luck that would undermine the quality of one’s knowledge but without depriving the subject of knowledge. The difference between knowledge-undermining luck and knowledge-preventing luck is a matter of degrees—the degree to which one’s entitlement-based explanation for the *veritic explanandum* is adequate. An unacceptably inadequate explanation would be a farfetched explanation and thereby lead to knowledge-preventing luck. However, there is no cut-off point distinguishing the two types of epistemic luck from each other—the notion of epistemic luck is seen as a spectrum concept for G-EL.

The second difference lies in that G-EL would not attempt to explain the knowledge-presence intuition in the fake-barn-county case by denying that Barney suffers from environmental luck. That is, as we have noted before, if

we apply the entitlement-based account of epistemic luck to the fake-barn-county case strictly, then Barney's entitlement propositions cannot adequately explain her epistemic success. After all, Barney barely enjoys any unearned epistemic entitlement due to her abnormal epistemological environment, while mere earned entitlements can hardly constitute an adequate explanation. The upshot is, Barney's belief must suffer from environmental luck—the problem is just how damaging this luck is. G-EL consents the luckiness of Barney's belief, and reconciles the disagreement between knowledge-presence intuition and the knowledge-absence intuition by diagnosing it as a disagreement on whether Barney's luck is knowledge-preventing. G-EL will admit that, the fake-barn-county case is a borderline case that falls at the vague boundary area between knowledge-undermining luck and knowledge-preventing luck. It is not a paradigmatic case for either type of luck. If one sees the fact that Barney happens to glimpse at the only real barn as an essential explanandum for explaining Barney's epistemic success, then Barney's entitlement propositions only provide a farfetched explanation as it fails to explain away an essential explanandum. If that is the case, then Barney's luck is knowledge-preventing. On the contrary, if one does not regard Barney's accidental glimpse at the real barn as an important explanandum related to her epistemic success, then Barney's luck is just knowledge-undermining. In that case, Barney's earned epistemic entitlements, *i.e.*, her justified beliefs such as 'I see a barn', would only fail to explain away a relatively unimportant explanandum. Therefore, the corresponding explanation for her epistemic success is inadequate (after all, the occasionality of her glimpse is an unaddressed relevant explanandum), but not farfetched.

In summary, for G-EL, the disagreement between the knowledge-presence intuition and the knowledge-absence intuition is a disagreement on the importance of the occasionality of Barney's looking at the real barn *qua* an explanandum. Both intuitions are reasonable when it comes to a borderline case like Barney's case. Moreover, there is no need to make a final verdict on

which side is correct, just like there is no need to make a final verdict on whether a shade of half-green-half-blue colour is green or blue. Green and blue are both spectrum concepts, just like knowledge-undermining luck and knowledge-preventing luck. G-EL can thus reconcile both intuitions.

So what about the FAKE RING COLLECTION series? For the original FAKE RING COLLECTION case, G-EL can argue that the relevant entitlement-based explanation would be farfetched as it is unexplained why the lady's friends happen to see her only genuine ring on the only genuine-ring-wearing date—this unexplained explanandum is essential for the corresponding belief's being true. While as the frequency of the lady's wearing fake rings steps down, the abnormality of her wearing genuine rings decreases as well. The upshot is, we find the unexplained explanandum, *viz*, why her friends happen to see her wearing a genuine ring, becomes less and less relevant for comprehending why her friends' beliefs are true. In particular, when it comes to FAKE RING COLLECTION III, it becomes counterintuitive to see her friends' epistemic conditions as *abnormal*. Consider Martin Smith's account of *normality*. According to Smith (2016), an abnormal occurrence calls for explanations, while a normal one does not. It is also arguable that likelihood is an important factor influencing the normality of an event⁷⁹. In the case of FAKE RING COLLECTION III, what requires an explanation is 'why does the lady wear her fake ring today' rather than why she does not. Thus, G-EL can account for the case of FAKE RING COLLECTION III by diagnosing that the relevant *veritic explanandum* can be *relatively* adequately explained in terms of the lady's friends' corresponding entitlement propositions. Thus it is unclear whether her friends suffer from *knowledge-undermining luck* (let alone knowledge-preventing luck) or not. Anyhow, the relevant entitlement-based explanation is not farfetched.

⁷⁹ Although Smith (2016) does not deny the connection between likelihood and normalcy, he holds that it is also normal for one to win the lottery despite the low probability of winning. I am not sure that I share the same intuition with Smith on this point. It strikes to me as plausible to claim that it is less normal to have a neighbour who won the lottery than to have one who never wins.

G-EL can also answer those problems raised in the FAKE BAR case. The Bar only sells surrogates on Sundays, thus it is certain that Dick and Harry would never have the chance to drink fake gins in the bar. We do not need to explain why they do not happen to imbibe the surrogates given the stipulations of the case—their drinking the surrogates would be *abnormal*. Hence, we are inclined to ascribe knowledge to Dick and Harry. On the contrary, since Tom goes out nearly every night, it is normal that he might have drunk the surrogate. In light of this, his corresponding entitlement-based explanation is less adequate—it strikes me that his corresponding explanation is as inadequate as Barney's. Thus, Tom's epistemic status falls at an intermediate range of scale that is borderline and equivocal.

So what about Oscar? Everyone tells Oscar that 'I'm drinking gin', does Oscar obtain testimonial knowledge from everyone? By G-EL, it should be fair to conclude that the luckiness of Tom's belief is transmitted to his testimony. Hence to account for the truth of Oscar's testimonial belief, it needs to be explained why Tom happens to visit the bar on a genuine-gin-selling day—Oscar's corresponding entitlement propositions can hardly accomplish this. In addition, if Oscar does not know when Dick, Harry and Tom usually appear in the Bar (or when they do not), and he just takes their testimonies at face value, then there would be at least one more unexplained explanandum for Oscar's relevant entitlement-based explanation. That is, why does Oscar happen to query Tom (rather than anybody else) on Tuesday (rather than any other day). This can be relevant to the explanation for Oscar's epistemic status, as if he asked a different customer on a different date, he could have easily formed a false belief via the same method. Therefore, it is less reasonable for us to ascribe knowledge to Oscar than to Tom regarding what Tom is drinking. Likewise, it is *less* intuitive (though *overall*, it might still be intuitive) that we should recognise Oscar as knowing what Dick and Harry are drinking compared with Dick and Harry themselves. To sum up, compared with Dick, Harry and Tom, the degree to which we are confident to ascribe

testimonial knowledge to Oscar should be downgraded in light of the extra unexplained explanandum.

In summary, G-EL's solution to the reconciliation problem differs from traditional absolutist approaches in that it does not aim to draw a threshold for knowledge or epistemic luck. For G-EL, what really matters is not *making verdicts* on Gettier-ed protagonists' epistemic statuses and then justifying these verdicts. Rather, G-EL focuses on *explaining disagreements* on Gettier-ed protagonists' epistemic statuses and then justifying these explanations. According to G-EL, most disagreements on whether a Gettier-ed subject knows are disagreements on whether those explananda that cannot be adequately explained by one's entitlement propositions are important. To be more specific, when assessing Barney's epistemic status in the fake-barn case, people disagree on whether Barney knows that there is a real barn, because people disagree on *to what degree* the fact that she happens to glimpse at the only real barn *is important*. The more important the unexplained fact is for an assessor such that it must be taken seriously, the more likely an assessor would decline to ascribe knowledge to Barney. On the contrary, for an assessor who finds the occasionality of Barney's glimpse unimportant or irrelevant, it is more intuitive that Barney possesses knowledge despite the unexplained environmental luck. The importance of the unexplained explananda, G-EL would be glad to admit that, is a matter of degrees and can often be subjective. There is no such a threshold for this kind of importance that can distinguish knowledge-undermining luck from knowledge-preventing luck. Also, G-EL can accommodate the intuition that, even for those who ascribe knowledge to Barney, Barney's knowledge is not a paradigmatic case of knowledge—after all, it is clear that Barney's belief at least suffers from knowledge-undermining luck. Hence, G-EL can solve the problem of reconciliation better without suffering from difficulties for both the steadfast and the sophisticated absolutist solutions.

Drawing inspiration from gradualism's successful application in addressing the problem of environmental luck, we can also rethink the strategy that we should take to solve the Gettier problem. Traditionally, the original Gettier problem and subsequent Gettier-style counterexamples are seen as challenges to various analyses to the concept of knowledge (such as the long-held JTB account, and subsequent JTB+X account). However, once a new JTB+X analysis or a new XTB analysis was put forward, it seems that there could always be new corresponding Gettier-ed counterexamples to the new analysis. With K-SPECTRUM in play, the threshold paradigm of epistemological investigation (see Chapter 5) that attempts to come up with a un-Gettier-able analysis of knowledge is conducting a strategical mistake. Instead, the spectrum paradigm of epistemological investigation will suggest that perhaps epistemologists' task is not to *avoid* Gettier-ed cases, but to *raise* more Gettier-ed cases and *explain* the intuitions behind those Gettier-ed cases. For a spectrum paradigm of epistemology, to obtain a more advanced and fine-grounded understanding of the concept of knowledge, what we need is not a set of necessary and sufficient conditions for knowledge, but more *conceptual indicators* of knowledge and non-knowledge. Non-knowledge indicators indicate the absence of knowledge. Some non-knowledge indicators are paradigmatic and uncontroversial, for example, falsity and disbelief. Some non-knowledge indicators are less salient but can be still seen as paradigmatic, for example, intervening luck. In contrast, there are conceptual indicators that are not paradigmatic and can invite controversy on whether they count as non-knowledge indicators—for example, environmental luck. The supplement of the list of paradigmatic non-knowledge indicators and our discussions of controversial potential non-knowledge indicators can both help us to boost our understandings of the extension of 'knowledge'. These two types of boost can be archived by putting forwards and debating over more Gettier-ed cases.

2. Gradualism and Faultless Disagreement

2.1: The Puzzle of Faultless Disagreement

Disagreements are ubiquitous. Ordinarily, when two agents disagree with each other, one would turn out to be at fault. However, sometimes, when two competent epistemic peers disagree with each other, and both have equally good epistemic support, we do not think that any one of them is making a mistake or at fault. Call this type of disagreement *faultless disagreement*. Faultless disagreements are often seen on matters of taste (e.g., on whether a sculpture is well-carved), but most people agree that there are also faultless disagreements on knowledge ascriptions⁸⁰. For example, many contextualists hold that the disagreement between an anti-sceptic and a sceptic (on, say, whether one knows that one has hands) is a paradigmatic case of faultless disagreement. They both are right in terms of their own epistemic standards respectively—the difference lies in that the standard is low for anti-sceptics, and high for sceptics. For readers who are not convinced by the contextualist solution to scepticism (like me), consider cases as follows:

Bank Case Keith and his wife are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Although they generally like to deposit their paychecks as soon as possible, it is not especially important in this case that they are deposited right away, so Keith suggests that they drive straight home and deposit their paychecks on Saturday morning. Keith's wife says, 'Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.' Keith replies, 'No, I *know* it'll be open. I was just there two weeks

⁸⁰ See Kölbel (2004) for an exception, where Kölbel presents an argument against the existence of faultless disagreements, while he also argues against that argument in the paper.

ago on Saturday. It's open until noon.' Keith's wife replies, 'Banks do change their hours. You *don't* know it'll be open tomorrow.'

This is a case modified from the well-known bank case from DeRose (1992: 913). Now in this modified case, Keith disagrees with his wife on whether he knows that the bank will be open tomorrow. Keith asserts that 'yes'—this seems to be correct, as he has good evidence (he went to the bank on Saturday two weeks ago), and he knows that they are in a low stake context to the effect that it does not really matter even if the bank does not open tomorrow. So his epistemic standard operative in this case is relatively low. Keith's wife asserts that 'no'—this seems to be correct as well, as she also has a non-conclusive defeater (banks do change their opening hours) and she knows that Keith fails to rule out the relevant error possibility. So her epistemic standard is relatively high. Their assertions are seemingly contradictory, but no one seems to make a mistake in accordance with their own epistemic standard.

There are many other cases of faultless disagreement on knowledge ascriptions. For example, disagreements on whether Barney knows in the fake-barn case, or whether Oscar knows in the fake-bar case. No matter what one's position is, it does not seem to be at fault.

Despite its widespread existence, the phenomenon of faultless disagreement is still puzzling. That is, when two agents disagree on the truth value of p , to wit, A asserts that p , and B asserts that not- p , it seems to be necessary that one agent's assertion must be *false*. If that is the case, then why none of them is *at fault*? This puzzle is in particular salient when it comes to faultless disagreements on knowledge ascriptions, as a knowledge ascription proposition is ordinarily taken to be of one and only one truth value. One cannot both know that p and fail to know that p . Call this puzzle the puzzle of *faultless disagreement*. The crux of this puzzle is to accommodate two

prima facie conflicting intuitions: when two agents A and B disagree on the truth value of a knowledge ascription proposition, 1) they are genuinely disagreeing with each other (call this intuition DISAGREEMENT); 2) none of them is genuinely making a mistake (call this intuition FAULTLESSNESS).

There has been intense debate regarding the solution to the puzzle of faultless disagreement. Eminent candidates include *indexical contextualism*, *non-indexical contextualism*, and *non-indexical relativism*. Before we proceed to how gradualism can help to solve the puzzle, let us analyse the pros and cons of those extant candidates.

2.2: Extant Proposals

*Indexical contextualism*⁸¹ is a view that the semantic contents of assertions depend on the context of utterance. It sees 'knows' as an indexical term like 'I' and 'today', of which the semantic meaning varies from context to context. Accordingly, the semantic contents of two assertions that constitute a faultless disagreement are incomplete. For example, in the bank case, what Keith really means is 'I know that the bank will be open tomorrow in accordance with *my* epistemic standard', and his wife actually conveys a message that 'You do not know that the bank will be open tomorrow in accordance with *my* epistemic standard'. In this way, the two seemingly incompatible assertions are both true in terms of their own epistemic standard. As long as one endorses the contextualist semantics of 'knows', indexical contextualism can apparently succeed in explaining FAULTLESSNESS.

⁸¹ Indexical contextualism is sometimes also known as 'indexical relativism' (see Kölbel 2003; Kompa 2015). It is 'indexical', because it consents that the contents of assertions are influenced. It is less easy to understand why this position is also seen to be 'relativist' by some philosophers. Here is my guess: Generally, the difference between contextualism and relativism lies in that the former claims that the truth-value of an assertion is determined by *the context of utterance*, while the latter focuses on *the context of assessment*. However, some times the distinction between the two types of context are too vague such that they seem to collapse into each other. For example, see Ian M. Church's critique to McKenna (2012), where Church charges that McKenna's indexical contextualism has no substantial difference with an indexical relativist position (Church 2012).

The problem lies in DISAGREEMENT. A common criticism to indexical contextualism is that it fails to explain why two agents are genuinely disagreeing with each other rather than just talking past each other, if they are just asserting that one knows or does not know in terms of their own standards (see MacFarlane 2007; Kölbel 2004). After all, 'By my epistemic standard A, S knows that p' and 'By my epistemic standard B, S does not know that p' are not genuinely incompatible. By indexical contextualism, Keith and his wife are not even disagreeing on the same proposition. Some possible replies to this criticism are available for indexical contextualists. For example, Davis (2007) argues that speakers are often 'semantically blind' such that they are not aware that their unuttered/omitted epistemic standards are the real divergence. López de Sa (2008) proposes a 'commonality presupposition' account for DISAGREEMENT. That is, in a case of faultless disagreement, to two speakers' assertions implies a 'commonality presupposition' that their epistemic standards are alike. Therefore they take themselves to be contradictory to each other. Neither replies are satisfactory, because they both admit that there is no genuine disagreement regarding the same proposition. For Davis, the two debating speakers *blindly* take their assertions to be incompatible due to their semantic blindness; for López de Sa, they *mistakenly* take their assertions to be incompatible because of the misleading commonality presupposition. Once the two debating speakers complete their assertions by indicating their personal epistemic standard, for example, by saying that 'you do not know, by 'know' I mean that ...', both Davis and López de Sa should agree that the two speakers would stop taking themselves to contradict each other. Their replies might be plausible from a contextualist perspective, but at the cost of abandoning DISAGREEMENT.

A more influential defence for indexical contextualism is Keith DeRose's 'single scoreboard semantics' (see DeRose 2004). DeRose argues that ordinary criticism to contextualism presumes that contextualism has to appeal to the semantics to the effect that there are multiple personal standards

determining the truth-condition of an assertion in a conversation. Inspired by the metaphor of 'scoreboard' in Lewis (1979), DeRose calls this *multiple scoreboards semantics*. By the multiple scoreboards semantics, in a conversation where faultless disagreement is involved, the two speakers assess knowledge ascription propositions by their personal scoreboards that do not overlap with each other. Consequently, they talk past each other. DeRose suggests that contextualists can avoid this issue by adopting a *single scoreboard semantics*, which argues that there is a shared single scoreboard in a given conversation that helps to determine the truth-condition of a proposition in question. Speakers can change the score by their conversational manoeuvres, but there is only a single score, at any given time, that governs the truth-conditions of all speakers' uses of 'knows'.

But it is far from clear that what determines or counts as a score at a given time. Apart from this ambiguity, MacFarlane (2007) objects that the single scoreboard semantics fails to give us enough disagreement. He points out that in addition to intra-conversational disagreements that happen between two conversing speakers (e.g., Keith disagreeing with his wife), there is also another important type of disagreement independent of conversations (e.g., self-disagreements, or disagreements between two scientist groups that do not know the existence of each other). MacFarlane charges that DeRose only explains intra-conversational disagreements, leaving more inter-conversational disagreements unaccounted for. The notion of 'scoreboard', by its inventor Lewis, is rooted in conversational language games. According to DeRose, the 'score' registered in the scoreboard is also determined by speakers' conversational manoeuvres. While in situations where there are no genuine conversations between the two sides of a disagreement, especially in cases where the two sides do not even know each other, it seems to be fair to conclude that the single scoreboard semantics is not applicable.

Instead, MacFarlane (2009) advocates a version of *non-indexical contextualism* (aka 'moderate relativism', see Kompa 2015; Recanati 2008).

Non-indexical contextualism holds that the truth-values of assertions vary with the contexts of utterance, but the contents do not. An utterance in a context C is true iff it is true as evaluated relative to the epistemic standard operative in the context of utterance. To put it more formally,

'NIC: An utterance of a sentence *s* (of the form 'A knows that *P*') in context *c* is true only if the proposition thereby expressed is true relative to circumstances (w_c, e_c), where w_c is the world of *c* and e_c is the epistemic standard operative in *c*.' (Kompa 2015: 143)

It differs from indexical contextualism in the sense that contexts only influence the truth-values, rather than the semantic contents, of assertions, and thus two disagreeing agents are disagreeing on the same proposition. However, by determining truth conditions of propositions by different epistemic standards, the same problem remains: why is there any genuine disagreement? As Kompa questions:

'We seem to disagree, yet on the assumption that truth is relative there is no longer any point in that. What you said may well be true relative to your standard while what I said may well be true relative to my standard. Who needs to disagree? If all disagreement were to be modelled along the nonindexicalist line, it would come out as mere perspectivalness, or so it seems.' (Kompa 2015: 147)

One may try to evade this criticism by arguing that we are actually disagreeing on which epistemic standard should be employed in this context of utterance. But this gives rise to another problem: Is there a single epistemic standard that best (or at least better) suits a given context *c* such that it should be the correct e_c operative in *c*? If non-indexical contextualists' answer is 'yes', then why it is not the case that the speaker who employs a wrong/improper epistemic standard that is making a mistake? For example, consider DeRose's original bank case. When Keith and his wife is in a high-stake epistemic context to the effect that they have to deposit their paychecks before Monday, it would be much more improper for Keith to insist that he

knows that the bank will be open on Saturday. His low epistemic standard is not suitable for the high-stake context, and thus he would be seen as at fault by ascribing knowledge to himself in terms of the low epistemic standard. Therefore, if there is a single correct/best epistemic standard for a given context, then non-indexical contextualism would fail to accommodate FAULTLESS. However, if non-indexical contextualism replies that there is no single best or correct e_c operative in c , to wit, there is more than one proper epistemic standard for a given context, then the difficulty in addressing DISAGREEMENT would be waiting on the corner⁸². In summary, non-indexical contextualism is subject to a dilemma: on the one hand, if there is a single correct epistemic standard that should be operative in a given context, then it will fail to explain FAULTLESS; on the other hand, if there is no such a single epistemic standard, then it will fail to explain DISAGREEMENT.

The third option is *non-indexical relativism* (aka, 'radical relativism', see Kompa, 2015), which is a view that the truth-value of assertions vary with different standards of assessment. It agrees with non-indexical contextualism on the point that contexts only affect the truth-value, rather than the semantic content, of an assertion. It differs from non-indexical contextualism in the sense that it appeals to the epistemic standard operative in the context of *assessment*, rather than the context of utterance. Accordingly, when two assessors A and B dispute about whether Keith knows in the bank case, non-indexical relativism would explain that A and B are both correct in making a

⁸² To address this issue, one might resort to the 'metalinguistic negotiation' account from Plunkett & Sundell (2013). The 'metalinguistic negotiation' account argues that when disagreeing on p 's truth value by employing different standards of 'knows', the two disagreeing speakers are actually conveying messages to the effect that 'knows' ought or ought not be used such that it applies to one speaker's assertion. Their conveyed message are contradictory. This account can hardly help non-indexical contextualism. First, it also suffers from the same defect of DeRose's single scoreboard semantics – it neglects non-conversational disagreements. Second, it also neglects situations where the two sides of a debate *take themselves to be* using 'knows' in the same way, *viz*, situations where the two speakers *hold* that they are employing the same epistemic standard, but are only disagreeing on whether the same standard is met. Last but not least, it says nothing about whether one side of the disagreement is correct. To be more specific, whether is there a better way that 'knows' *ought to be* used in the given context? If yes, then why FAULTLESS? If no, then what is the point of the disagreement? Why it is not the case that both speakers are *mistakenly* insisting that their own epistemic standards are *better*?

judgment in terms of their own epistemic standards that operative in their own contexts of assessment (rather than Keith's context of utterance). The divergence lies in their different standards of assessment. It seems to be able to avoid the defect of being limited to conversational disagreements, as no matter whether there is a conversation, when a disagreement occurs, it would always be a disagreement between two assessors who assess the truth value of a proposition.

However, non-indexical relativism also faces widespread criticisms. Apart from general objections to the relativist semantics of 'knows' (see Boghossian 2006; Zimmerman 2006; Cappelen and Hawthorne 2009; Seidel 2014; Carter 2016, Ch. 7; *etc.*), Kompa (2015) also worries that non-indexical relativism is incompatible with the widely-accepted 'norm of assertion': one should only assert what is true. If relativism is right in that the truth-value of an assertion is determined by the context of assessment, then how can an utterer know what epistemic standards would assessors employ in different contexts of assessment? And, how can an utterer know what is true before (s)he asserts anything? If one cannot know all of these, then how can one obey the norm of assertion? Moreover, non-indexical relativism also seems to suffer from a similar dilemma that non-indexical contextualism is subject to. In many cases of faultless disagreement, the two sides of assessment seem to be in the same context of assessment. For example, Keith and his wife in the bank case, or two persons disputing about whether Barney knows in the fake barn case. Now the question is: Is there a single correct/best epistemic standard that should be operative in a given context of utterance? If yes, then why FAULTLESS? If no, then why DISAGREEMENT? Transferring the focus from contexts of utterance to contexts of assessment seems to do no help to solve this dilemma.

2.3: The Gradualist Solution

When disagreeing with others on knowledge ascriptions, e.g., 'Barney knows that there is a barn in front of her', we are disagreeing on the truth value of a given knowledge ascription proposition. Given the absolutist tradition, it is apparently intuitive that Barney either knows or does not know that there is a barn, to wit, the knowledge ascription proposition must be either true or false. The truth-condition of the proposition in question seems to be subject to the threshold of knowledge—if Barney's epistemic status meets the threshold, then the knowledge-ascribing side wins; if it does not, then the knowledge-denial side wins. Accordingly, the disagreement, in an absolutist picture, is a disagreement on the threshold of knowledge—whether the controversial case in question falls into the left or the right side of the threshold? Also, there can be two ways to interpret the relevant disagreement: 1) The two sides take themselves to be using different thresholds, and thereby disagree on which threshold should be applied to the given context⁸³; 2) Both sides take themselves to be using the same threshold, and thus disagree on the assessment of details of Barney's epistemic status. No matter which interpretation is endorsed, the aforementioned dilemma would follow: If there is a single correct threshold, then why FAULTLESS? After all, with the threshold in play, one side of the disagreement must have made a *false* verdict. If there is no such a threshold, then what is the point of DISAGREEMENT? The two sides can be compatible with each other by clarifying their own threshold.

Under an absolutist framework, when accommodating FAULTLESS, extant proposals introduced before would find it difficult to reconcile K-THRESHOLD and FAULTLESS. There seems to be an intrinsic confliction between the two theses. If T_k exists and whether T_k is met is the essential divergence of the disagreement, then one side of the disagreement will perforce be at fault—either be at fault for using the wrong T_k , or for using the correct T_k in a wrong way so that a false verdict is derived. Call this conundrum *the problem of*

⁸³ As we have seen in Chapter 4, contextualism can also be compatible with K-THRESHOLD by holding that for any given context, there is a corresponding threshold for knowledge.

faulty verdict. While in explaining DISAGREEMENT, extant proposals would encounter difficulty in reconciling *impurist* K-THRESHOLD and the existence of genuine contradiction. To evade *the problem of faulty verdict*, contextualists or relativists are inclined to conclude that both sides of the disagreement are making a correct verdict within their own context using their own threshold. This gives rise to several defects: first, it fails to explain cases where both sides are in the same context using similar standards; second, if they are in two parallel contexts using two distinctly different standards, then they seem to be just talking past each other. Call this conundrum *the problem of parallel contexts*.

A gradualist solution to the puzzle of faultless disagreement can avoid the two problems by replacing K-THRESHOLD with K-SPECTRUM. Consider situations where two persons disagree on whether a shade of colour among the spectrum is a colour of orange or red. The shade of colour in question falls right into the intermediate range between orange and red, so it is reasonable for both sides of the disagreement to argue for their own judgments. However, they are actually disagreeing with each other on the same topic (*i.e.*, what colour it is) in the same context using two highly similar standards. They are not using the words 'orange' or 'red' in two distinct ways—they might understand the *intensions* of the two words in the same way, they would reach quick agreement on what is a paradigmatic instance of red (or orange)—they are just debating on the *extensions* of the two types of colour in a borderline case. Similarly, consider situations where two customers in a steakhouse arguing about whether a cut of steak is medium-rare or rare. The doneness of the steak might just be between medium-rare and rare—if might be half-medium-rare and half-rare, so it is far from clear which degree of doneness the steak is of. In this case, none of the two customers is at fault for making their own verdicts, but they are genuinely disagreeing with each other by defending two contradictory verdicts.

Analogously, if knowledge is recognised as a spectrum concept, then cases of faultless disagreement on knowledge ascriptions can be seen as disagreement on borderline cases that falls into the intermediate range between knowledge and non-knowledge. When disagreeing on whether Barney has knowledge in the fake barn case, people are genuinely disagreeing with each other on the very *same* knowledge ascription proposition ‘Barney knows that there is a barn in front of her’. Gradualism does not need to adopt the indexical position that the semantic contents of the proposition in question are incomplete. Also, it does not need to appeal to different contexts of utterance or assessment to explain why neither side is making a mistake. To accommodate DISAGREEMENT, gradualism can explain that there is a genuine disagreement regarding whether the case under debate is a borderline case of knowledge or a borderline case of non-knowledge. The two sides of the disagreement are in the same context of assessment, sharing highly similar understandings of what knowledge is, just like agents share similar understandings of what red is in the orange/red and customers share similar understandings of what a medium-rare steak typically be like. They are disagreeing on the vague boundary of the extension of ‘knows’, just like when people disagreeing on whether a half-orange-half-red apple is a red apple or an orange apple.

Therefore, gradualism can avoid *the problem of parallel contexts*—it is in the same context that the disagreement on the same proposition occurs. A genuine disagreement does not need to be hinged on ‘threshold’—there can be disagreements focusing on the intermediate range of a spectrum. Faultless disagreements on knowledge ascriptions are this type of disagreement.

So how does gradualism accommodate FAULTLESS if the two speakers are disagreeing on the truth value of the same knowledge ascription proposition? How can gradualism avoid *the problem of faulty verdict*? Is there a single correct verdict to the truth value of the controversial knowledge ascription?

With K-SPECTRUM in play, gradualism’s reply is: no. In fact, when it comes to the intermediate range and borderline cases, there is no non-arbitrary criterion defining the truth-condition of a controversial knowledge ascription proposition. That is because, the truth-condition of ‘S knows that *p*’ depends on what the extensions of ‘knows’ are—while this is exactly what the two sides of the dispute are disagreeing on. There is not an objective and neutral third-party criterion that serves as the threshold helping us to determine the truth-conditions of propositions in controversy. We do not come to form the conception of knowledge by appealing to a presumptive threshold. The (vague) boundary of knowledge’s extensions is shaped by controversies and debates. When we reach controversies where faultless disagreements emerge, we know that we reach the boundary of knowledge. Therefore, it is natural that there would not be a presumptive single correct verdict for a sufficiently controversial borderline case. Just like there ought not to be a presumptive standard answer to whether a shade of borderline colour is red or orange. To be more specific, consider DeRose’s account of ‘gappy truth-condition’:

‘Suppose the person you’re talking with makes a statement with ‘gappy’ truth-conditions. Take, for example, a standard case of vagueness. She says ‘Frank is tall,’ where her claim is true if Frank’s height is in a certain range, is false where Frank’s height is in a certain lower range, and, we’ll suppose, is neither true nor false if Frank’s height is in an intermediate range.’ (DeRose 2004: 16)

Accordingly, DeRose introduces a trivalent system for assessing truth values of knowledge ascription propositions:

<u>P</u>	<u>It is true that P</u>
T	T
F	F
N	F

In the trivalent truth table above, ‘N’ refers to a proposition’s being gappy, *i.e.*, neither true nor false. MacFarlane (2007) grants that this gap view can accommodate FAULTLESS properly (his main discontent is with DeRose’s explanation of DISAGREEMENT), but I worry that this trivalent system cannot be applied in an absolutist picture. That is because, with the threshold for knowledge in play, a controversial knowledge ascription proposition is either true or false, to wit, one either meets T_k or does not—this is an absolute yes-or-no affair (recall the core thesis of intra-threshold absolutism that we noted in Chapter 3). There is not a third ‘gappy’ option. In contrast, within a gradualist framework, we can adopt a trivalent system by drawing on DeRose’s gap view:

<u>P</u>	<u>It is true that P</u>	<u>not-P</u>	<u>It is true that not-P</u>
T	T	F	F
F	F	T	T
B	B	B	B

In this modified truth table, ‘B’ means being borderline. A proposition p ’s truth value is B, *iff*, either (1) p conveys a knowledge ascription that makes a verdict one’s epistemic status in a borderline case that falls into the intermediate range in the spectrum of knowledge; or (2) p conveys a high-order proposition evaluating the truth value of a knowledge ascription proposition whose truth value is B.

With this trivalent system in play, neither party of a faultless disagreement is making a *false* verdict, and this system is compatible with K-SPECTRUM but not K-THRESHOLD. Gradualism can thus explain FAULTLESS without giving rise to *the problem of fault verdict*.

In summary, a gradualist solution can better solve the puzzle of faultless disagreement by avoiding *the problem of parallel contexts* and *the problem of fault verdict* that trouble an absolutist solution. Another advantage of this

gradualist solution is that it does not need to be based on a controversial semantics foundation such as relativist semantics or contextualist semantics, and hence is less theory-laden. Therefore, gradualism can better accommodate both FAULTLESS and DISAGREEMENT without suffering from defects of extant proposals within an absolutist picture.

3. Concluding Remarks

So far, we have seen that gradualism enjoys stronger explanatory power when addressing the asymmetry problem, the reconciliation problem of environmental luck, and the puzzle of faultless disagreement. These advantages discussed might not be able to nail down the conclusion that gradualism is a better reading of knowledge than absolutism. But I hope that I have said enough to enable the readers to conceive the promising prospect of developing the spectrum paradigm of epistemological investigation, and to consider abandoning the equivocal attitude on the gradability problem of knowledge. The final chapter of this thesis will continue to explore gradualism's application in solving another (perhaps the most) crucial epistemological issue: the sceptical problem.

Chapter Six: Epistemic Gradualism and Scepticism

Abstract: This chapter aims to explore another theoretical merit of epistemic gradualism. I will demonstrate that if one is to solve the sceptical problem, then gradualism is preferable to absolutism. Influenced by the absolutist orthodoxy, we find the sceptical problem paradoxical because we are committed to a claim that one's belief that 'I am not a brain-in-a-vat' does not meet the threshold for knowledge. An absolutist anti-sceptical approach rejects this claim by proving that there is a threshold that is met. I will prove that this absolutist approach would give rise to two problems: 1) absolutists can hardly locate such a threshold, because if there is a threshold for knowledge then it ought to be both less and more lenient than safety/sensitivity; 2) it risks begging the question *qua* an anti-sceptical approach. In contrast, gradualism denies the existence of knowledge's threshold and thereby blocks the sceptical argument in a more undercutting manner. The conclusion is that gradualism enjoys advantages over absolutism, at least in terms of its application in solving the sceptical paradox.

As we have noted before, one essential agenda for the debate between absolutism and gradualism is around the truth of K-THRESHOLD. Insofar as one grants gradualism, one should strive to reject or at least undermine K-THRESHOLD. This is not an easy job for gradualists, given that talk of knowledge's boundary is deeply entrenched in the orthodox epistemological discussion. Nonetheless, there are several potential argumentative strategies available for gradualists. For example, if it can be shown that rejecting K-THRESHOLD would be helpful to solve important epistemological issues, then gradualism would be credited with good reasons to overthrow external absolutism. In this chapter, I will explore one possible argument for gradualism that takes this strategy. I will demonstrate how gradualism can be

applied to address the sceptical problem by rejecting K-THRESHOLD. I do not think that what my argument below will provide a robust reply to the sceptical problem. Rather, my focus is not solving the sceptical paradox, but showing why gradualism is preferable to absolutism when solving the sceptical paradox. By doing so, I hope to enable the reader to judge the plausibility of the general argumentative strategy that gradualism can adopt to eliminate the equivocal attitude towards the gradualism/absolutism debate.

1. The Sceptical Paradox

Radical scepticism purports to show that we do not possess any everyday knowledge that we take ourselves to possess. It is widely accepted that the sceptical problem, by its nature, is presenting a paradox for us to solve (see Pritchard 2009b & 2015b; Wright 2008; Byrne 2004; *etc.*). The sceptical paradox is constituted by a set of intuitively appealing while *incompatible* propositions that we are committed to, and aims to reveal the inconsistency within our fundamental epistemological commitments. For example, the most-discussed *closure-based sceptical paradox* runs as follows:

[Closure-Based Radical Sceptical Paradox]

(PC1) If one knows that p and that p entails q , then one knows that q .

(PC2) I do not know that I am not a brain-in-a-vat (hereafter, BIV).

(PC3) I have widespread everyday knowledge.

PC3 is one of our most basic epistemological commitments. Rejecting PC3 simply leads to an epistemic hazard. PC1 is a simple formulation of the closure principle, which is referred to as 'known entailment closure' (see Bernecker 2012: 368). Of course, there is more than one formulation of the closure principle (see Hawthorne 2004; Klein 2004; Dretske 2005; Pritchard 2015b). However, for the sake of simplicity, focusing on this simple form of

closure will be enough for the purpose of this chapter⁸⁴. Surely, there is a long-held contention over what the most reasonable formulation of closure is, but it is much less contentious that closure enjoys widespread successful applications in our daily epistemic practice. I know that I am sitting in front of my laptop, and I know that sitting is not standing, and thus by closure I know that I am not standing. In addition, closure is also essential for extending our knowledge by deduction from what we already know⁸⁵. Given the remarkable epistemological importance of closure, closure-based scepticism also becomes ‘the most popular way of motivating radical scepticism in the contemporary literature’ (Schönbaumsfeld 2016: 7). There are anti-sceptical proposals in the literature that attempts to reject closure. For example, Dretske (1970) and Nozick (1981) famously argue that closure should be abandoned due to its incompatibility with the sensibility condition of knowledge⁸⁶. However, apart from extensive criticism of the sensitivity principle⁸⁷, abandoning closure is also deemed as an unaffordable epistemological cost by many epistemologists (see Pritchard 2015b; Ashton 2015; *etc.*). In addition, many have argued that the sceptical paradox can be constructed independently of the closure principle—for instance, the underdetermination principle and the indistinguishability principle can both serve as the replacement (see Vogel 1997; Schönbaumsfeld 2017; Pritchard 2015b). In consideration of these, denying PC1 might not be a satisfactory resolution to the sceptical paradox. In contrast, no matter what epistemic principle sceptics employ to impose conditions to knowledge, they would always need a claim that those conditions are not met by our everyday beliefs—that is exactly the role that PC2 is supposed to play in the sceptical argument. Given that neither rejecting PC1 nor discarding PC3 is an

⁸⁴ The reader is welcomed to substitute other reasonable forms of closure into PC1 – this would not affect my main argument, as I do not plan to put weight on closure to solve the sceptical paradox.

⁸⁵ This idea is typically reflected by a sophisticated formulation of closure, *i.e.*, *the competent deduction closure*: if a subject S knows that *p* and comes to believe that *q* by completely deducing *q* from *p*, then S knows that *q*.

⁸⁶ Hetherington (2001; 2011) also proposes to reject closure. We will discuss his anti-closure arguments latter.

⁸⁷ For example, see Vogel (1987), Sosa (1999), and Pritchard (2008).

appropriate way to solve the sceptical paradox, there remains PC2. So why do we find PC2 appealing?

PC2 involves a *prima facie* appealing sceptical hypothesis. The sceptical scenario stipulates that a BIV's illusory perceptual experiences are introspectively indistinguishable from those of her non-envatted duplicate. Moreover, if one were a BIV, one would still believe the opposite. Thus, it seems that one is not able to distinguish a BIV-world from a not-BIV-world. To further motivate PC2, sceptics would ordinarily appeal to a certain epistemic principle (e.g., sensitivity, underdetermination, indistinguishability) to the effect that one's belief of 'I am not a BIV' falls short of knowledge. For an external absolutist who advocates K-THRESHOLD, being committed to PC2 implies being committed to the following thesis:

UNMET: My belief that 'I am not a BIV' does not meet T_k .

As we have noted in Chapter 4, T_k can be formulated in terms of different relevant scales that can be used to evaluate one's epistemic status. What is the most relevant scale that we should focus on in the present discussion? The answer relies on what sceptics actually *doubt* when they motivate UNMET. Ordinarily, sceptics do not doubt that you are genuinely *believing* that you are not a BIV. They could also concede that it might be *true* that you are not a BIV—the standard sceptical reasoning only doubt whether you *know* this to be true, given that you do not have stronger epistemic support enabling you to rule out the sceptical scenarios. They doubt that your evidence can ever sufficiently justify your belief; they doubt that you have better reasons to support that \sim BIV rather than BIV; they doubt that you can discriminate the BIV-worlds from those non-BIV ones... What sceptics *typically* doubt by putting forward UNMET, so to speak, is *the strength of epistemic support* that your beliefs enjoy. Hence the form of T_k that I will be discussing in this chapter is one concerning the strength of epistemic support. Failing to reach this type of threshold for knowledge, our belief of the

denial of sceptical hypothesis, combined with the closure principle, would entail the sceptical conclusion by *modus tollens*. Consequently, influenced by the absolutist tradition, the sceptical paradox is apparently paradoxical for us, since the following three epistemological commitments of us are inconsistent:

- (1) **CLOSURE:** If one knows that p and that p entails q , then one knows that q .
- (2) **UNMET:** My belief that 'I am not a BIV' does not meet T_k .
- (3) **EK:** I have widespread everyday knowledge.

Among the three claims, EK might be the most fundamental epistemological commitment underpinning our rational system. The truth of EK is ordinarily taken for granted when it does not encounter the challenge from the sceptical argument. The sceptical paradox brings a huge impact on our epistemological conception because it threatens the otherwise indubitable foundation of our rational system, *i.e.*, EK. The closure-based sceptical argument threatens the truth of EK by employing CLOSURE and UNMET. It has been argued that abandoning CLOSURE is not an ideal choice, consequently, rejecting UNMET would be a natural anti-sceptical manoeuvre. If UNMET can be disproved, then there is no reason to deny that EK obtains.

But wait! Cannot the sceptical paradox be constructed without appealing to the concept of 'threshold for knowledge', but by employing mere *necessary conditions* of knowledge? It seems that sceptics just need to impose a certain *necessary condition* C_k to knowledge (*e.g.*, sensitivity or distinguishability) that one's belief of the denial of the sceptical hypothesis fails to meet. Hence, there seems to be another apparently plausible diagnosis of the sceptical paradox:

- (1*) **CLOSURE:** If one knows that p and that p entails q , then one knows that q .

(2*) **UNMET***: My belief that 'I am not a BIV' does not meet C_k .

(3*) **EK**: I have widespread everyday knowledge.

It looks like that UNMET* is obviously distinct from UNMET. However, the *prima facie* difference between the two claims could be revealed to be merely verbal and superficial by a more intensive analysis. As we have noted before, EK is our most fundamental epistemological commitment that would otherwise be taken for granted if the sceptical argument fails to provide sufficient reasons for us to doubt its truth. For a sceptical argument employing UNMET*, it is UNMET* that provides the *ground of doubt* as to EK's truth. If UNMET* is falsified, then scepticism would fail to imperil EK. In another word, if one's belief that 'I am not a BIV' is proven to be able to satisfy C_k , then there would be no reason to refuse to credit her with the relevant knowledge. Hence, it is plausible to conclude that, for whoever that attempts to get rid of scepticism, not-UNMET* entails EK. Likewise, as we have noted before, EK can also be entailed from not-UNMET for the same reason. Meanwhile, the sceptical paradox indicates that UNMET entails not-EK, and UNMET* also entails not-EK. By contraposition, EK entails not-UNMET, and EK entails not-UNMET*. Thus, we have 'not-UNMET* \rightarrow EK', and 'EK \rightarrow not-UNMET'. Also, 'not-UNMET \rightarrow EK', and 'EK \rightarrow not-UNMET*'. Combining these four conditionals, we have 'UNMET \leftrightarrow UNMET*'.

One might worry that: the derivation of the biconditional above is derived too quickly. After all, sceptics might argue that even if C_k is satisfied, there might be another higher condition for knowledge that we failed to meet. Thus one cannot conclude that not-UNMET* entails EK. In this case, the onus is on sceptics to put forward a new C_k to reconstruct the sceptical argument. But this does no help to reject our conclusion that once (the new) C_k is satisfied, there is no reason to ascribe knowledge. Sceptics can iterate this manoeuvre by keeping raising the epistemic standard delivered by C_k , but this would only make the C_k employed in UNMET* more and more akin to T_k *per se*. Think

about what we (*qua* anti-sceptics) should do if sceptics keep raising the epistemic standard conveyed by C_k . If we fail to satisfy the relatively lower-standard C_{k0} put forward by sceptics, then of course our everyday knowledge would be doomed. But if we prove that a low-standard C_{k0} can be satisfied, then either sceptics have to admit that knowledge follows from the failure of UNMET*, or they could employ a higher-standard C_{k1} to reformulate UNMET*. In any case, for this run-and-catch game, it is unchanged that once the currently highest C_{kn} at issue is met, then before a higher C_{kn+1} is advocated, everyday knowledge should be ascribed to us. Eventually, it would be the highest C_k that sceptics can reasonably⁸⁸ employ that serves as the last barrier needing us to overcome. But now, what is the difference between this highest reasonable C_k and the threshold for knowledge? They both refer to a cut-off point where a piece of belief starts being counted as knowledge, a boundary between knowledge and not-knowledge.

This can also explain why UNMET* and UNMET have no substantial argumentative difference. Remember that a threshold typically comes in the form of a necessary and sufficient condition. In UNMET*, C_k is claimed to be a mere necessary condition of knowledge. Sceptics use C_k as a barrier hindering knowledge from being obtained. Nonetheless, in light of EK's great appeal, once C_k is satisfied, the barrier is hurdled, then knowledge would follow. This indicates that C_k is virtually *constituting* a not only necessary but also, *ceteris paribus*, sufficient condition for a true belief's being knowledge—at least, it is (provisionally) sufficient before sceptics put forward a more demanding C_k . For every given sceptical argument, the point where a piece of belief satisfies C_k , *ceteris paribus*, is exactly the point where we have no reason to refuse to recognise a piece of true belief as knowledge, to wit, the point where the piece of true belief starts counting as knowledge—that is exactly the definition of a 'threshold'. For every given formulation of UNMET*, the corresponding C_k delimits an epistemic condition such that once this

⁸⁸ If sceptics formulate C_k in an unreasonably over-demanding manner, then it would fail to constitute a paradox. We would not find UNMET* (so formulated) appealing at the first place if it involves an excessively high C_k .

condition is satisfied, then here comes knowledge, otherwise, ignorance—it divides beliefs that are qualified as knowledge from beliefs that are not qualified enough. An epistemic condition like this, is virtually a threshold for knowledge—it plays the exact role of T_k , *viz*, distinguishing knowledge from everything that falls short of knowledge. That is to say, before sceptics reformulate the sceptical argument by imposing a $C_{k_{n+1}}$ that is more demanding while still reasonable, every given C_k employed by a given version of the sceptical argument is argumentatively equivalent to T_k . Surely, sceptics can try to formulate C_k in different ways and thereby construct different forms of sceptical argument. If the C_k employed is too permissive, then UNMET* can hardly follow; while if the C_k employed is too demanding, then it would hardly be seen as a reasonable necessary condition for knowledge which can make the sceptical argument appealing. Therefore, eventually, the best-constructed sceptical argument should involve the optimal formulation of C_k , which is the most demanding C_k that still remains plausible. And this optimal formulation of C_k for sceptics, so described, is also identical with the most plausible formulation of T_k that one can ever conceive. In summary, for every given sceptical argument that appeals to UNMET*, C_k is dialectically identical with T_k ; for the best-constructed (and thus most challenging and philosophically significant) sceptical argument, the best formulation of C_k that sceptics can utilise is identical with the most plausible formulation of T_k that absolutists can put forward. Hence, the difference between UNMET* and UNMET is far from being clear and substantial for both absolutists and gradualists. Being committed to UNMET* is dialectically tantamount to being committed to UNMET.

With these being said, it should be fair to diagnose the paradoxical nature of scepticism as hinging on our commitment to UNMET. Refuting UNMET would be an ideal choice for resolving the sceptical paradox. Precisely put, UNMET can be spelt out as: there is a threshold for knowledge T_k , which my belief that 'I am not a BIV' fails to meet. Accordingly, there are two ways to reject UNMET:

Approach 1: Admitting the existence of T_k , and proving that our belief that 'I am not a BIV' meets T_k .

Approach 2: Denying the existence of T_k .

Given the distinction between K-THRESHOLD and K-SPECTRUM, it is not hard to see that Approach 1 is an absolutist approach, while Approach 2 is a gradualist one. In what follows, I will prove that Approach 2 is preferable to Approach 1. Therefore, rejecting external absolutism can provide us with a better solution to the sceptical paradox.

2. Threshold and Scepticism

I will develop two arguments to show why Approach 1 is not a preferable option for rejecting UNMET. The first argument will examine, if we grant that there is a T_k that can both falsify UNMET and serve as a reasonable necessary and sufficient condition of knowledge, then what this sort of T_k could be. The idea is, if we can find a candidate of T_k as a reference, and demonstrate that 'if such a T_k exists, then it should be both more permissive and less permissive than the candidate', then by *reductio ad absurdum*, we should deny the existence of such a T_k . Call this *the argument from reference*. The second argument will reveal the intrinsic dialectical defect of Approach 1 as an anti-sceptical manoeuvre. That is, by presupposing the existence of knowledge's threshold, Approach 1, which typically comes in the form of a neo-Moorean anti-sceptical proposal, will inevitably presume the existence of knowledge. This kind of anti-sceptical manoeuvre would be begging the question. Call this *the argument from petitio principia*.

Let me start with the argument from reference. I will first choose two plausible candidate thresholds, namely, safety and sensitivity, as the references. A dilemma for absolutists will be revealed which shows that a satisfactory anti-

sceptical T_k is impossible. After that, I will address two potential worries to this argument: that we ought to exam more plausible scales/dimensions, and that my argument is just a simple repeat of the Gettier problem.

2.1: The Argument from Reference

Safety and Sensitivity

The Island of Pianosa is recruiting pilots. The recruitment requirements state that the candidates' height should be not less than 6 feet, but less than the height of the Air Force General Dreedle. Question: what is the threshold for being a pilot in the Island of Pianosa? If it turns out that General Dreedle is exactly 6-foot tall, then we can conclude that such a threshold does not exist: no one can be both less and not less than 6-foot tall. The threshold itself is incoherent. We do not even need to try any other potential candidate of the threshold. We do not need to ask: What about 6.1 feet? What about 5.8 feet? Would 1.89 metres fit the bill? That is because we have already known that there is one *reference*—namely, 6 feet—that suffices to indicate the absurdity of the threshold.

Analogously, we can consider the issue regarding knowledge's threshold in the same vein. It is undeniably hard for absolutists to draw a clear threshold for knowledge, and they can reasonably argue that they do not need to find the exact location of the cut-off point distinguishing knowledge from ignorance in order for them to know the existence of such a cut-off point (this is the attitude of moderate K-THRESHOLD). Therefore, a charitable objection to absolutists should allow them to just provide some acceptable approximations. We could start figuring out the approximated T_k by appealing to some reasonable references. If it turns out that we find there is a reference T_r such that T_k should be both more permissive and less permissive than T_r , then we can conclude that T_k does not exist, without needing to exhaustively

examine all possible candidates for T_k . In that case, T_r is to T_k what '6 feet' is to the regulated height of pivots in Island Pianosa.

Assume that we are going to take approach 1 to reject UNMET. In that case, the essential task is to find out a threshold for knowledge (namely, T_k), which satisfies two conditions: 1) it is met by our belief that 'I am not a BIV', such that UNMET can be falsified; 2) it is qualified to provide us with a proper necessary and sufficient condition of knowledge, such that it can serve as a reasonable threshold. What would such a T_k be like? To narrow down the range where the threshold is located, we can choose a hopeful approximation as a reference. Name this reference T_r . We can analysis T_r by inspecting whether it satisfies the two conditions that T_k is expected to meet. If it is shown that T_r is too lenient to fit the bills, then we can conclude that T_k ought to be more demanding than T_r . On the contrary, if it turns out that T_r is unacceptably strict, then T_k should be more permissive than T_r . So what reference can we utilise?

T_k can be defined in terms of various scales that are relevant to knowledge, so it can be formulated in numerous different ways. Which scale should we care about here? Now that we are examining the absolutist reply to scepticism, absolutists are expected to (at least approximately) provide a threshold for knowledge that can reject UNMET. By advocating UNMET, sceptics are questioning the strength of epistemic support for our beliefs that \sim BIV. Accordingly, absolutists should provide a T_k that can answer the question of *how strong* the epistemic support for a belief should be in order for that belief's being knowledge. One important way to evaluate this kind of strength that epistemologists nowadays usually consider is to examine a belief's *fallibility* or *luckiness*⁸⁹. Epistemologists such as Hetherington believe that the less fallible one's knowledge is, the stronger one's epistemic support is. How fallible is one's knowledge allowed to be? When does one's true

⁸⁹ For discussions of the connection between fallibility and epistemic luck, see Hiller & Neta (2007); Booth (2010).

belief start being too fallible to be counted as knowledge? A popular answer is: when it is merely luckily true. Pritchard (2007) famously defines epistemic luck in terms of *safety*: a belief is luckily true when it is unsafe. He delicately defines safety as follows:

S's belief is safe *iff* in most near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true. (Pritchard 2007: 292)

So it seems that *safe true belief* (abbr. 'SaTB') will be a natural candidate of T_r that we can take into consideration. Another motivation for considering SaTB is its dialectical connection with UNMET. A classic sceptical reasoning supporting UNMET requires that in order for one to know \sim BIV, one has to be able to discriminate \sim BIV scenarios from BIV scenarios. Another way to put it is to require one's belief that \sim BIV to be *sensitive*: that is, if one were a BIV, one would stop believing that one is not a BIV via the same method M that one uses in the actual \sim BIV world. Since one's belief that \sim BIV does not meet this requirement, UNMET obtains. The safety-based neo-Mooreanism famously attempts to block this sceptical reasoning by resorting to the safety condition for knowledge. The idea is, one's true belief that one is not a victim of a sceptical hypothesis could be recognised as knowledge by virtue of its being safe—it does not need to be sensitive:

'One of the key issues facing neo-Mooreanism is how to explain how we can know the denials of radical sceptical hypotheses. The standard line in this regard usually adverts to some form of the safety principle for knowledge, as defended, for example, by Ernest Sosa (1999). This principle holds, roughly, that what is essential to knowledge is that one has a belief that could not have easily been false. The basic idea is that provided sceptical error-possibilities are indeed farfetched, then it follows that one's true belief that one is not a victim of such an error-possibility will be such

that it couldn't have easily been false, and so can count as knowledge.' (Pritchard 2012b: 116)

Hence it is natural to take SaTB as a considerable candidate of T_k that can falsify UNMET, and thereby PC2, so that the sceptical argument could be disarmed. Safety and sensitivity both evaluate the quality of knowledge along the scale of the strength of epistemic support (and fallibility as well).

Admittedly, absolutists do not *have to* formulate T_k merely along this scale. But now that absolutists are to offer a robust anti-sceptical proposal on the basis of T_k , an ideal formulation of T_k should be able to meet the challenge that sceptics raise from perspectives where safety and sensitivity are relevant.

With this being said, now let us examine SaTB *qua* a reference (T_r) for drawing T_k . Remember that an ideal T_k should meet two conditions: being able to falsify UNMET and being able to serve as a plausible necessary and sufficient condition for knowledge. By examining whether T_r (in this case, SaTB) meets the two conditions of T_k , we will be able to figure out what an approximation of T_k should be like such that it can meet both conditions. The idea is, if we find that T_r fails to meet one or two of the condition(s), then we should ask what causes these failures. And then those reasons of failures can be seen as the defects of T_r that our T_k should ultimately avoid. However, if we find there are some failures that are inevitable—just like the 6-foot-dilemma that we have seen at the start of this section—then it should be concluded that an ultimately satisfactory T_k is impossible. Let us start by examining whether SaTB, *qua* a T_r , meets the first condition.

According to the safety condition of knowledge, one's belief that 'I am not a BIV' would not have easily been false, given that in most near-by and all very close near-by possible worlds where one continues to believe so with the same method (e.g., 'seeing that I have both hands and then inferring that I am not a BIV'), it won't be the case that one is a BIV indeed. There are at least two ways to understand why: first, a neo-Moorean can argue that every

sceptical scenario world is an outlandish world, rather than a near-by world; second, in a world where one is a BIV, one would not be able to use the exact *same method* to form her relevant belief, as a BIV does not actually *see that* she has both hands (insofar as we endorse that 'seeing-that-*p* implies *p*'). In conclusion, one's belief that 'I am not a BIV' is not only true but also *safely true*, and hence ought to count as knowledge. It seems that safety, *qua* an example of T_r , can ideally protect our knowledge of the denial of sceptical hypotheses from being deprived by UNMET. So the first condition for a T_r to be a proper T_k is satisfied. But what about the second condition? Does safety constitute an acceptable threshold for knowledge? Is 'true and safe belief' a proper necessary and sufficient condition of knowledge?

Many would say 'no'. Some argue that SaTB is *too permissive* for defining knowledge, as there are true beliefs that are safe but not knowledge.

Consider E. J. Coffman's PRESUMPTUOUS SECRETARY⁹⁰:

PRESUMPTUOUS SECRETARY: An eminent historian, Hank, recently discovered that Abraham Lincoln was born not in 1809 (as most of us think) but in 1806. Hank writes a letter to his friend, Sandy, in which he asserts that Lincoln was born in 1806. When preparing the letter to be sent, Hank's presumptuous new secretary assumes he made a careless mistake about Lincoln's date of birth, and changes the text so that it says Lincoln was born in 1809. Unbeknownst to all, the secretary's printer has just developed the following glitch: when directed to print a '9', some other numeral besides '9' is randomly selected and printed instead. As luck would have it, Hank's letter gets printed as stating that Lincoln was born in 1806. When

⁹⁰ Another type cases of safe non-knowledge involve knowledge of necessary truths. Since in no possible world would a necessarily true proposition (e.g., mathematical truths) be false, one's true belief of these propositions will never violate the safety condition of knowledge. However, one can still form a true belief of necessary truths by virtue of mere luck, for example, one can luckily form a true belief that ' $12 \times 13 = 156$ ' by trusting a broken calculator that accidentally generated the correct answer. Pritchard (2009, 2012a) address this issue by extending the account of safety to the effect that it covers similar beliefs that one forms via the same method in similar circumstances, rather than just being limited on the single particular belief that the subject forms in the actual world. However, it is unclear what counts as a similar belief and a similar circumstance.

Sandy receives the letter, she consults it for the answer to the question of when Lincoln was born. Sandy comes to believe, justifiedly and truly, that Lincoln was born in 1806. (Coffman 2010: 246)

Intuitively, Sandy does not know that Lincoln was born in 1806, as there is something that lucky significantly undermines her knowledge. However, her true belief is safe, as while reading Hank's letter, she could not have easily formed (via the same method, namely, reading Hank's letter) a false belief. The letter's text has been settled and is now stable; also, Sandy's cognitive faculties are working properly. So in most nearby possible worlds where Sandy continues to form the belief that Lincoln was born in 1806 by reading Hank's letter, her belief would continue to be true. Her belief is a case of safe non-knowledge.

Grundmann (2018) objects to this case by arguing that although Sandy's belief is certainly Gettier-ed, it is not safe. Grundmann describes Sandy's method of forming beliefs as 'believing whatever he reads in Hank's letter', which is an unsafe method that could have easily generated false beliefs:

'But now consider the relevant modal environment. In some nearby possible worlds the printer is directed to print "9" but randomly prints some other numeral, e.g. "8" or "7". So it might easily have printed "1808" or "1807". In these cases, Sandy would have acquired a false belief about when Lincoln was born, namely that Lincoln was born in 1808 or that he was born in 1807.' (2018: 8)

Goldberg (2015) also reads Coffman's case in a similar manner. He remarks that:

'Coffman's example succeeds against SM only on a somewhat contentious construal of the method involved in cases of relying on written testimony. If the method is characterized in such a way that it does not initiate with the written report itself, but instead begins with those processes that produced the written product, those processes (it would seem) could easily have produced a false belief—as would

have been the case had there been no glitch in the printer.’
(2015: 277)

I think neither Grundmann nor Goldberg substantially undermines Coffman’s case. It is not hard for Coffman to slightly modify his case such that Sandy forms her belief via a less fishy method. For example, the printer’s glitch can be described as “when directed to print a ‘9’, the numeral ‘6’ will be randomly selected and printed instead”⁹¹. More importantly, even if we keep the original stipulated descriptions in the case of PRESUMPTUOUS SECRETARY, Grundmann and Goldberg fail to show that Sandy’s belief in the original case is unsafe in accordance with Pritchard’s delicate definition of safety. According to Pritchard’s definition of safety, in most nearby and all very close possible worlds where Sandy continues to form a belief of *the target proposition* (in this case, ‘Lincoln was born in 1806’) with the same method (in this case—presumptively Grundmann and Goldberg are right—‘believing whatever Hank wrote in the letter’—and we read this method as initiating with those processes that produced the written product), her belief (*i.e.*, ‘Lincoln was born in 1806’) would still be true. So this satisfies Pritchard’s classic definition of ‘safety’. Grundmann and Goldberg at most show that Sandy’s belief-forming *method* is unsafe—in the sense that it could have easily produced beliefs of some false propositions other than the target proposition—but the genuine central problem here is whether Sandy’s *belief* of the target proposition is safe. They fail to show that Sandy’s *belief* is unsafe.

⁹¹ Some may then argue that in this revised case, the intuitive verdict is Sandy *knows* that Lincoln was born in 1807. I do not find this verdict intuitive. Consider Pritchard’s case of TEMP (see Pritchard 2012a). In the case of TEMP, Temp forms true beliefs about the temperature in the room by consulting a thermometer, which is in fact broken and controlled by a hidden agent who ensures that every time Temp consults the thermometer, the reading is always correct. It is widely accepted that Temp lacks of knowledge. And I do not see substantial difference between the way in which Temp’s beliefs are lucky and the way in which Sandy’s belief is lucky -- the former is by virtue of a hidden agent, while the latter is by virtue of an odd glitch. They both ensure that the protagonists’ beliefs are safely true, and it does not strike that it really matters whether it is a living agent or a machine’s glitch.

So, if SaTB is too lenient for drawing an ideal threshold for knowledge, then we should expect T_k to be less permissive than SaTB. One way is to impose more conditions to knowledge such that we can raise the threshold for the strength of epistemic support. For example, in order to account for cases such as TEMP (another case where one's belief is safely true but fails to be knowledge, see Pritchard 2012a) and DEMON (a case where one's belief is not only safely true but also manifests one's cognitive competence but still fails to be knowledge, see Pritchard 2015a), Pritchard (2015a) argues that the simple conjunction of a true belief's *being safe* and *manifesting the cognitive competence* is not sufficient to guarantee knowledge if the former has nothing to do with the latter. Instead, he suggests that we should also require the safety of one's true belief is sufficiently attributable to the manifestation of one's cognitive competence. This is also the gist of his anti-luck virtue epistemology. Does this constitute an ideal T_k this time? I'm afraid that the answer is still 'no'⁹². One important reason that I want to emphasise here is: there are also reasonable grounds for us to think that safety is not necessary at all for knowledge. Thus in whatever way anti-luck virtue epistemologists formulate T_k , it will always contain an unnecessary condition imposed to knowledge, namely, safety, that makes the proposed T_k too demanding. If that is the case, then absolutists will face a dilemma: on the one hand, T_k should be less permissive than SaTB; on the other hand, T_k should be more permissive than SaTB.

SaTB is also taken to be *not* lenient enough, as some argue that the safety conditions would preclude some intuitively appealing cases of (unsafe) knowledge. For example, Chris Kelp (2009: 20) raises a case against the safety condition of knowledge:

ARCH-NEMESIS: Suppose Russell's arch-nemesis has an interest that Russell forms a belief (no matter

⁹² For other criticisms of anti-luck virtue epistemology, see Kelp (2013), Carter (2013), Broncano-Berrocal (2014), Zhao (2018), etc.

whether true or not) that it's 8:22 by looking at the grandfather clock when he comes down the stairs. Russell's arch-nemesis is prepared to do whatever it may take in order to ensure that Russell acquires a belief that it's 8:22 by looking at the grandfather clock when he comes down the stairs. (Since we are concerned with a conceptual claim here, Russell's arch-nemesis may have means available to do so that we can imagine only in our wildest dreams. For instance, he may be an evil-demon who can set the clock to 8:22 with his invisible hand a second before Russell looks at it.) However, Russell's arch-nemesis is also lazy. He will act only if Russell does not come down the stairs at 8:22 of his own accord. Suppose, as it so happens, Russell does come down the stairs at 8:22. Russell's arch-nemesis remains inactive. Russell forms a belief that it's 8:22. It is 8:22. The grandfather clock is working reliably as always.⁹³

Kelp argues that, albeit unsafe, Russell's belief that 'it is 8:22 now' still counts as knowledge. Bogardus (2012) argues that Kelp's counterexample does not successfully present us with a case of unsafe knowledge. The idea, succinctly put, is that Russell was in epistemic danger before he decided to walk down the stairs at 8:22, but once he decided to come down the stairs at 8:22, the danger dissipated and his belief was formed safely. For Bogardus, *S was at epistemic risk just before forming her belief* does not imply that *S's belief was formed unsafely*. The real counterexample to the safety condition in Bogardus's mind goes as follows:

ATOMIC CLOCK The world's most accurate atomic clock hangs in Smith's office. The clock is accurate by virtue of

⁹³ For a similar case, see Bogardus (2012).

its sensitive radiation sensor, which, however, could easily malfunction if a radioactive isotope were to decay in the vicinity. Today, unbeknownst to Smith, a radioactive isotope was accidentally left on the table near-by the clock. The isotope could decay at any moment, but luckily, it didn't start to decay when Smith entered the office and looked at the normally running clock that read '8:22 am'. Smith thus formed the true belief that he entered the office at 8:22 am. (See Bogardus 2012: 300-301.)

Bogardus holds that, Smith's belief is obviously unsafe even after he decided to enter the office at 8:22, as the radioactive isotope still could have easily decayed and thus the clock's reading could have easily been false. Nevertheless, it is intuitive that Smith still knew that he entered the office at 8:22 am. The lack of safety does not entail the lack of knowledge. Bogardus diagnoses that 'safety at a time depends counterfactually on what would likely happen at that time or soon after in a way that knowledge does not. That, ultimately, is why knowledge need not be safe.' (2012: 305). Bogardus's ATOMIC CLOCK was challenged by Broncano-Berrocal (2014). Broncano-Berrocal argues that Smith's belief is actually safe, because in a near-by world where the clock malfunction, Smith would be forming the belief via a different method. Broncano-Berrocal's objection is essentially rest on a highly questionable distinction: in the actual world, the method m_1 that Smith uses to form the belief is 'looking at the working clock'; while in the near-by possible world, the method m_2 that Smith uses is 'looking at the stopped clock' (see Broncano-Berrocal 2014: 77). The two methods, so defined, are surely different. However, this way of characterising belief-forming *method* looks *ad hoc* and excludes many paradigmatic cases of unsafe belief, for example, Russell's original STOPPED CLOCK case, and Goldman's FAKE BARN COUNTY case. Take the latter as example, in accordance with Broncano-Berrocal's taxonomy, Barney's belief is also safe, because if she happens to glimpse at a fake barn, her belief-forming method would be 'looking at a fake

barn', differing from that in the actual world, *i.e.*, 'looking at a real barn'. In light of this, Broncano-Berrocal's requirement for 'same method' is unreasonably demanding. For a full reply to Broncano-Berrocal's objection, see Bogardus & Marxen (2013).

Now it seems that Broncano-Berrocal's objection to Bogardus's ATOMIC CLOCK is not very successful. At the meantime, one might find Bogardus's criticism of Kelp's ARCH-NEMESIS kind of odd. After all, in ARCH-NEMESIS, Russell could have easily decided not to walk down at 8:22 (alternatively, he might not make any relevant *decision* on purpose at all, but just randomly come down to the stairs at a random time), why cannot these close possible worlds threaten the safety of his belief? If they can, then it seems like both Kelp's and Bogardus's cases serve as cases of unsafe knowledge.

However, proponents of safety have another more steadfast response to the two cases. Pritchard (2009b) replies to Kelp's case by claiming that he does not share the intuition that Russell really have knowledge. He argues that Russell's true belief is just a cognitive achievement that falls short of knowledge because of its unsafety. For Pritchard, the distinction between cognitive achievement and knowledge is reflected clearly in the case of FAKE BARN COUNTY. That is, Barney's cognitive success is sufficiently attributable to her cognitive agency and thus counts as her cognitive achievement. Nevertheless, Barney lacks knowledge because her belief suffers from environmental luck. Kelp also admits that any version of safety that renders Barney's belief unsafe would also render Russell's belief unsafe. So now the crux seems to lie at whether Barney *knows* or not. This is a knotty issue that we have discussed in the previous chapter. The underlying question here is whether FAKE-BARN-COUNTY-style environmental luck is incompatible with knowledge or not. As we have noted, there are two conflicting intuitions regarding whether Barney lacks knowledge in the fake barn county. I will keep my stance on this issue neutral (as I did in Chapter 5), and turn to show that if one is to defend SaTB by this steadfast response,

then this will perforce yield another dilemma for absolutists—there will be a new T_r that is both too lenient and too demanding indicating the impossibility of T_k . This time, the new T_r will be one that is closely connected to SaTB, namely, *sensitive true belief* (SeTB).

Before we look more closely at how SeTB can be involved in our argument, let us take stock of what has been achieved. We have seen that, in order to provide a robust reply to scepticism by taking Approach 1, absolutists are expected to find a T_k that can meet sceptics' challenge from the perspective of the strength of epistemic support for one's belief that \sim BIV. So fallibility/luckiness will be a relevant scale in this context, and SaTB is a natural choice of T_r , *i.e.*, a reference that helps us to carve out an ultimately satisfactory T_k . With SaTB in play, we find a dilemma for absolutists. For the first horn, if there is a T_k , then it should be less permissive than SaTB (given cases of safe non-knowledge such as PRESUMPTUOUS SECRETARY and TEMP). For the second horn, if there is a T_k , then it should also be more permissive than SaTB (given cases of unsafe knowledge such as ARCH-NEMESIS and ATOMIC CLOCK). One steadfast way to evade the second horn is to insist that in those putative cases of unsafe knowledge, the protagonists *lack* knowledge for the same reason that Barney lacks knowledge in FAKE BARN COUNTY. Therefore, for those who do not share the intuition that Barney lacks knowledge, this response does not hold water, and SaTB will suffice to show that it is impossible to have a T_k that is both less permissive and more permissive than SaTB. For those who share the intuition that knowledge is absent in FAKE-BARN-COUNTY-style cases, now I invite them to consider *sensitive true belief* as an alternative reference T_r .

When seeing SeTB as a new T_r , we will expect it to play the same role that SaTB did in our previous arguments. That is, if it turns out that a satisfactory T_k ought to be both less permissive and more permissive than SeTB, then we can conclude that T_k does not exist. The sensitivity condition of knowledge can be formulated as follows:

A subject S's true belief p is sensitive, iff, *in the closest possible worlds* where p is false, S would not believe that p via the same method M (see Nozick 1981; Pritchard 2008; *etc.*).

Sensitivity is famously advocated by Dretske (1970) and Nozick (1981). It is argued that, to constitute knowledge, a belief has to be not only true, but also sensitive. However, many have demonstrated that the sensitivity-based account of knowledge is not sufficient. Accordingly, by employing sensitivity as a reference for locating T_k , we would have: if there is a T_k , then it should be *less lenient* than sensitivity. Why is sensitivity too lenient? Consider Saul Kripke's RED BARN:

RED BARN: Rob visited a fake barn county full of papier-mâché barns that are indistinguishable from genuine barns. There is only one genuine barn in the county, which is red. Unbeknownst to Rob, every fake barn is green. Rob happened to look at the only real barn and thus formed the true belief that 'there is a red barn in front of me'. (see Kripke 2011: 186)

Did Rob know that there was a red (real) barn in front of him? If one holds that Barney lacks knowledge in the original FAKE BARN COUNTY case, one can hardly consistently credit Rob with knowledge in this RED BARN case. After all, Rob suffered from the same environmental luck as Barney does. If Barney does not know that 'there is a barn', so does not Rob. If Rob does not know that 'there is a barn', he also does not know that 'there is a red barn'. But was Rob's belief sensitive? Yes. That is because, Rob formed his belief by looking at the object and judging its colour. In most near-by possible worlds where Rob happened to see a papier-mâché barn instead, he would not believe that 'there is a *red* barn', as all fake barns are green. Thus, Rob's

true belief meets the requirement of sensitivity, but still fails to count as knowledge.

Likewise, many cases showing safety's insufficiency can also be used to reveal that sensitivity is too permissive. For example, it can be argued that in TEMP and DEMON, both protagonists have a sensitive true belief. Nevertheless, since sensitivity (as well as safety) fails to capture the sense that knowledge is a type of cognitive achievement, it is *too lenient* to be a proper threshold for knowledge.

On another hand, there is also a common worry that sensitivity might be *too demanding* a condition for a belief's being knowledge. A straightforward concern is that sensitivity implies PC2 and thus supports UNMET. One's belief that 'I am not a BIV' is insensitive as if one were a BIV, one would still believe that she is not regardless. In light of this, many have employed sensitivity to construct the sceptical argument (see DeRose 1995; Lehrer 2000; Sosa 2004; Leite 2004; Immerman 2016). So sensitivity, *qua* a T_r , does not meet the first condition of a proper T_k , namely, being able to falsify UNMET. Apart from this, Sosa's TRASH BAG case is also widely accepted as a counterexample towards the sensitivity condition:

TRASH BAG: On his way to the elevator Ernie releases a trash bag down the chute from his high rise condo. Presumably he knows the bag will soon be in the basement. But what if, having been released, the trash bag had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance? None such could affect Ernie's predictive belief as he releases it, so he would still predict that the bag would soon arrive in the basement. (see Sosa 1999: 145)

Arguably, Ernie knows that the bag will soon be in the basement, despite the insensitivity of his belief. Therefore, we have cases indicating that sensitivity is unnecessary for knowledge—sensitivity is too demanding. The upshot is, if there is a T_k , then it should be *more lenient* than sensitivity.

Now we have seen the whole picture of my argument from reference. It consists of two parts. The first part of my argument shows that, by using SaTB as a reference for T_k , we have:

- (AFR 1): If there is a T_k , then it should be less permissive than SaTB, as there are cases of safe true beliefs that do not qualify as knowledge.
- (AFR 2): If there is a T_k , then it should be more permissive than SaTB, as there are cases of unsafe knowledge.
- (AFR 3): It is impossible for a threshold to be located at a point that is both higher and lower than SaTB along the scale of the strength of epistemic support.
- (AFR 4): By *modus tollens*, the previous three premises show that the antecedent of AFR 1 and 2 should be rejected, that is to say, there is no T_k .

The second part of my argument considers a steadfast objection to AFR 2 insisting that as long as one's true belief is influenced by the environmental epistemic luck, one would always lack knowledge. I have argued that, if one is committed to this steadfast response, then one should also endorse that the protagonist in the case of RED FAKE BARN lacks knowledge for the same reason. Hence, the second part of my argument employs SeTB as a new reference T_r , which yields a new dilemma:

- (AFR 5): If there is a T_k , then it should be less permissive than SeTB, as there are cases of sensitive true beliefs that do not qualify as knowledge.

(AFR 6): If there is a T_k , then it should be more permissive than SeTB, as there are cases of insensitive knowledge.

(AFR 7): It is impossible for a threshold to be located at a point that is both higher and lower than SeTB along the scale of the strength of epistemic support.

(AFR 8): By *modus tollens*, the previous three premises show that the antecedent of AFR 5 and 6 should be rejected, that is to say, there is no T_k .

Undeniably, one might find SaTB and SeTB unattractive as candidates of T_k . I do not deny that there are some seemingly more promising candidates other than SaTB and SeTB. However, notice that, if my argument above succeeds, then we do not even need to consider any other candidate of T_k , as SaTB and SeTB would already suffice to reveal the incoherence within the idea of there exists a T_k . Just like we do not need to ask whether 6.1 feet or 6.13 feet or 5.98 feet would be a proper threshold for the height of pilots in Island Pianosa. To comprehend why, let us take a closer look at some possible objections to my argument.

Multi-Dimension Worry

One might argue that this only shows that SaTB is not a suitable candidate of T_k , rather than that T_k does not exist. For example, not all men can be employed as pilots in the USA; also, some female candidates can be employed as pilots in the USA. Hence, so to speak, 'being male' is both too strong and too weak for the threshold for being a pilot. But this does not indicate that the threshold for being a pilot in the USA does not exist. Instead, this just shows that 'being male' is neither a necessary nor a sufficient condition. However, this is an improper analogy. In the USA-pilot case, gender is not a relevant scale for being a pilot—it is not something that we care about for the purpose of recruiting excellent pilots. However, when it

comes to the case of anti-scepticism, as we have noted, the strength of epistemic support is one essential relevant scale for finding a T_k that can reject UNMET. SaTB (as well as SeTB) falls at this essential scale and yields a dilemma to the effect that no T_k that it is both less permissive and more permissive than SaTB can be located along this scale. A more proper analogy should be pilot recruitment in the Island Pianosa, where the height of candidates is stipulated as an essential scale for drawing the threshold for recruiting pilots. 'Six-foot-tall' falls at this essential scale, and yields a dilemma to the effect that no threshold can be both more lenient and less lenient than 'six-foot-tall'. SaTB is to knowledge what 'six-foot-tall' is to Island Pianosa's pilots.

The worry that we did not exhaust all possibilities of T_k for absolutists might be rooted in an idea as follows:

Even if we accept your readings of those Gettier-cases, what they would seem to show is just that any T_k must be less permissive than SaTB/SeTB *in some respects* and that any T_k must be more permissive than SaTB/SeTB *in some respects* from which it doesn't follow that there is no T_k . Obviously, there couldn't be a necessary and sufficient condition for knowledge that is both *logically stronger* and *logically weaker* than safety—but there could perfectly well be a logically *independent* condition that is weaker than safety in some ways and stronger in others. You seem to assume that any other candidate condition must be located on a one-dimensional scale where safety/sensitivity falls at. However, T_k can consist of various conditions that are located on multi-dimensional scales.

Name this *the multi-dimension worry*. As I have stated, T_k can be formulated in terms of various relevant scales, to wit, be constructed in a multi-

dimensional manner. What I aim to show is that the scale where SaTB and SeTB are located on is *an essential scale that absolutists cannot evade* when we are examining the absolutist anti-sceptical approach. That is to say, no matter how many other dimensions/scales that absolutists might appeal to in order to put forward a satisfactory T_k , the scale where SaTB and SeTB are involved is one that absolutists have to accommodate if they are to reject UNMET. I interpreted this scale as one concerning the strength of epistemic support. This scale/dimension of T_k asks: How strong does one's epistemic support need to be in order to yield knowledge? At the beginning of this section, we have seen that this question is very important for absolutist's finding a proper T_k because one's strength of epistemic support is exactly what UNMET doubts about. Another way to ask this question that Hetherington usually uses (when he talks about 'the boundary problem of knowledge') is: How fallible is one's true belief allowed to be? SaTB/SeTB answers that: 'To the extent that it is safe/sensitive'.

Still, one might further question that fallibility/luckiness is not the only dimension for absolutists to analyse 'the strength of epistemic support', and thus SaTB/SeTB is not the only way to answer this question about the strength of epistemic support. Of course, there can be many other ways to answer this question, but the answer offered by SaTB/SeTB has already indicated that *none of them* will succeed in providing us with a T_k that can address this question. Just like there are many ways to answer the question: How strong is a wrestler required to be in order to be signed by WWE? One might answer by 'being able to powerslam the Big Show', 'being able to deadlift 600 pounds', 'being stronger than Steve Austin', and so on. However, none of them can be true if it is known that 'any wrestler who wants to work for WWE has to be both more muscular and less muscular than John Cena'. Surely, 'being more/less muscular than John Cena' is just *one* way to answer the question; 'muscularity' is just *one* dimension to evaluate how strong an athlete is. But as long as it is true that 'any wrestler who wants to work for WWE has to be both more muscular and less muscular than John Cena',

then any other candidate answer/dimension would be meaningless—it has been shown that no matter how strong one is, one can never work for WWE, as no one can be both more and less muscular than John Cena. It is the same with the strength of epistemic support. SaTB/SeTB serves as an answer just like ‘being as muscular as John Cena’. As long as we acknowledge fallibility/luckiness as *one* proper dimension to evaluate one’s strength of epistemic support (among many other proper dimensions)—just like muscularity is one proper dimension to evaluate how strong a wrestler is—then the impossibility of T_k revealed by SaTB/SeTB also applies to any other candidate answer. So why is fallibility/luckiness a proper dimension? One reason, as we have noted before, is the dialectical connection between this dimension and UNMET. Another reason that I take to be plausible is: if absolutists are to put forward an *ultimately satisfactory* threshold for knowledge, then this ultimate threshold should capture the prevalent intuition that knowledge cannot merely be a result of luck. Of course, this intuition is not agreed by all epistemologists (Hetherington is exactly such an exception), but it should at least serve as a *conditional* motivation for acknowledging the dimension concerning safety/sensitivity: if you accept that knowledge is incompatible with epistemic luck, then a robust formulation of T_k should be able to capture this intuition. And the antecedent, I believe, will obtain for many people.

To put it formally, these are my claims:

Let the scale where SaTB and SeTB are located on be S_0 .
Let the scale of the strength of epistemic support be S_{sep} . I leave open whether S_{sep} is identical with S_0 ; I leave open whether S_{sep} consists of more dimensions apart from S_0 . My point is that S_0 is a both dialectically and conceptually *important* dimension for analysing S_{sep} — S_0 is one *proper* dimension (just like other proper dimensions, if any) that one can appeal to when evaluating the strength of epistemic

support. I also concede that it might be possible that ‘a putative T_k is logically stronger than SaTB/SeTB along a scale S_1 , and logically weaker than SaTB/SeTB along another scale S_2 , as long as S_1 and S_2 are two different scales, and they are not identical with S_{sep} or S_0 ’. It cannot be the case that ‘a putative T_k is both logically stronger and weaker than SaTB/SeTB along the scale S_{sep} or S_0 ’.

Absolutists propose to reject UNMET by coming up with an anti-sceptical T_k . SaTB and SeTB have shown that it is impossible for any anti-sceptical T_k to hold water on the scale S_0 , and thus S_{sep} . Given that S_{sep} is an essential scale that an anti-sceptical T_k has to accommodate, it is impossible for any T_k to hold water regardless of any other scale S_n .

With these being said, now it should be clear why my argument is not threatened by the multi-dimension worry. This also, again, explains why we do not need to examine more candidate forms of T_k .⁹⁴ Also notice that if my argument from reference succeeds, then it can be seen as not only criticism of the absolutist anti-sceptical approach but also a direct objection to K-THRESHOLD. The idea is, T_k cannot exist, as nothing can be both less permissive and more permissive than SaTB/SeTB. Hence we shall not see knowledge as a threshold concept.

The Gettier Problem

⁹⁴ To hammer home this point, we can try to consider some other analyses of knowledge that might be deemed more plausible than SaTB or SeTB, for example, Duncan Pritchard’s anti-luck virtue epistemology (2012a), or Chris Kelp’s knowledge first virtue epistemology (2017, 2018). Both analyses absorb Sosa’s AAA account of knowledge as an addition to other elements (for the former, the anti-luck conditions; for the latter, the knowledge-first spirit). Neither analyses can evade the scale which SaTB and SeTB are located on. This is obvious for Pritchard’s anti-luck virtue epistemology, as it explicitly includes the anti-luck condition for knowledge. It is the same with Kelp’s knowledge first virtue epistemology. That is because, being based on Sosa’s virtue epistemology model, it also needs to answer what a ‘competent belief’ (in a terminology that Sosa uses more frequently, ‘adroit belief’) is, which also addressed the problem regarding the strength of epistemic support. Besides, it is also noteworthy that Sosa’s own account of ‘adroitness’ adopts a reliabilist reading, which also concerns fallibility.

Now it seems like my argument from reference is just simply a repeat of the point that absolutists need to respond to the Gettier problem, rather than an objection to the absolutist anti-sceptical approach. In fact, I will admit that if one is to adopt Approach 1 to argue against scepticism, then one is subject to the Gettier problem. One standard way to solve the Gettier problem is to figure out what is the non-Gettier factor X that complements the traditional JTB tripartite analysis of knowledge such that we can give knowledge a proper set of necessary and jointly sufficient conditions. Notice just how similar a work this is to absolutists' task when locating a threshold for knowledge. If absolutists are able to provide a T_k that meets the two anti-sceptical conditions that we emphasise in this chapter, then this would thereby enable absolutist to give the Gettier problem an ideal solution in a sense. Correspondingly, if the Gettier problem is left unsolved due to absolutists' inability to find the necessary and sufficient conditions for knowledge, then absolutists would also fail to defend their anti-sceptical approach. That is the price of choosing Approach 1. The real key problem that we should ask is: Why won't the gradualist anti-sceptical proposal be plagued by the Gettier problem as well? That is because gradualists never promise to refute sceptics by defending a T_k . Quite the contrary, they never admit the existence of such a T_k . So they do not need to solve the sceptical problem or the Gettier problem by looking for the necessary and sufficient conditions for knowledge. Therefore, even though the Gettier problem *might* vex gradualists as well for some other independent reasons (remember that, however, in Chapter 5 we have also seen how gradualism can help us to do justice to the Gettier problem—the Gettier problem can be resolved without needing to find a T_k), it does not vex gradualists when gradualism is applied to solve the sceptical problem. In solving the sceptical problem by taking Approach 2, gradualists are not expected to also accommodate the Gettier problem.

Another question is, does this mean that if absolutists could come up with a condition that solves the Gettier problem, then they would also be able to respond to my argument from reference? No. I admit that if they succeed in finding an ultimately satisfactory T_k , then they might be able to also give the Gettier problem one solution by finding the X factor. But not vice versa. That is because there is more than one approach to solving the Gettier problem—other than coming up with a T_k . This opens up room for absolutists to solve the Gettier problem without thereby disarming my argument from reference. For example, absolutists can put knowledge first⁹⁵ and give the Gettier problem a Williamsonian reply—but this job has nothing to do with articulating what T_k is. Also, an absolutist can appeal to Sosa’s AAA account of knowledge to address the Gettier problem. In this case, it is unclear how this approach can respond to my argument from reference—the idea of ‘adroitness’ and ‘aptness’ would still need to answer the question regarding the strength of epistemic support, and thus would also encounter the dilemma revealed by SaTB/SeTB. The apt belief account of knowledge *might be* sufficient for addressing the Gettier problem (I am not claiming that it actually *is*), but it is not fine-grained enough for offering us a T_k that can block my argument from reference.

I hope to have said enough to show that the argument from reference is not a simple repeat of the Gettier problem. In contrast, it is an independent challenge to the absolutist anti-sceptical approach.

2.2: The Argument from *Petitio Principia*

One might find my argument from reference essentially hinge on one’s intuition about whether a case counts as a counterexample to safety or sensitivity. Is there anything else that gradualists can do if one holds an intuition contradicting mine? I believe that there is a more intrinsic defect of

⁹⁵ In this case, I think a knowledge-first absolutist might need to accept *moderate* K-THRESHOLD, as they give up analysing knowledge in terms of necessary and sufficient conditions. The idea here might be, we do not know where the threshold is, but there must be one.

anti-sceptical proposals that take Approach 1—they are risking of begging the question.

Anti-sceptical proposals embracing Approach 1 typically come in the form of neo-Mooreanism. Neo-Mooreans attempt to reject PC2 by arguing that a satisfactory epistemic standard that can best capture the nature of knowledge should be like such and such, and then proving that, in accordance with this epistemic standard, we *know* that we are not victims of sceptical hypotheses. Now the question is: can they propose such a reasonable epistemic standard without presuming our possession of widespread everyday knowledge? Take one of the most influential neo-Moorean proposals, *viz*, Sosa's safety-based proposal, as an example. How did Sosa motivate safety and show its advantage over sensitivity? The main method that Sosa (1999) took was to illustrate counterexamples to sensitivity that safety can properly avoid, *e.g.*, the TRASH BAG case. This method can work, significantly by virtue of our *epistemic intuition* that: according to our ordinary understanding of what knowledge is, Ernie can be seen as knowing that the bag will soon be in the basement—otherwise, many intuitive cases of knowledge would be excluded from the market of knowledge as well, which is an unpalatable upshot. Safety's plausibility is based on our epistemic intuition that safety can better reflect our ordinary understanding of knowledge. And this epistemic intuition is essentially grounded in our possession of widespread everyday knowledge. One finds safety more appealing than sensitivity, because the former can save knowledge that we ordinarily take ourselves to have (*e.g.*, predictive knowledge, as TRASH BAG shows) from being denied by sensitivity. If one is not committed to our possession of widespread everyday knowledge, how can epistemic intuition as such be grounded? With these being said, it should be more evident why a neo-Moorean anti-sceptical proposal that takes Approach 1 succumbs to the problem of *petitio principia*: If a neo-Moorean epistemic standard is at all plausible to us, then it must accord with our understandings of knowledge that we think we possess—that is the foundation of its plausibility. However, what the sceptical paradox imperils is

exactly this foundation. Hence, it is unclear how neo-Mooreanism can refute scepticism in a non-circular manner.

It is noteworthy that there are also some philosophers holding that the neo-Moorean reasoning is not begging the question. A prominent example is James Pryor (2004, 2012), whose view is also echoed by philosophers such as Allan Hazlett (2014). Pryor's argument is based on an account of perceptual justification that he calls 'dogmatism', as opposed to the 'conservative' view that can be found in Wright (1985), Cohen (2002), and White (2006), *etc.* Dogmatism denies that having *prima facie* perceptual justification to believe that 'here are two hands' requires one to have antecedent justification to believe that \sim BIV. Hence it is not the case that Moore's proof of the external world succumbs to the transmission failure of warrant. One can acquire justification to believe that there is an external world by having experiences of hands. Moore's proof is not begging the question as having justification to believe Moore's conclusion *is not among the conditions* that make one has the justification that one purports to have for the premises. Indeed, there is a sort of epistemic dependence between Moore's conclusion and the premises, but Pryor argues that it is not a vicious one.

I do not plan to dive into the ongoing debate between dogmatism and conservatism as this will be far beyond our purpose here. However, there are two points that I would like to emphasise.

First, the form of neo-Moorean reasoning that Pryor discussed was based on Moore's proof, which was formulated as follows:

- (M1) Here are two hands.
- (M2) If hands exist, then there is an external world.
- (M3) So there is an external world.

This is manifestly distinct from the form of neo-Moorean argument that I am analysing here:

(N1) I know that I am not a BIV, *because* my relevant belief meets T_k .

(N2) If I know that I am not a BIV, then the sceptical argument fails.

(N3) So, I have widespread everyday knowledge.

Or, alternatively:

(N1*) I know that I have hands, *because* my relevant belief meets T_k .

(N2*) If I know that I have hands, then I know that I am not a BIV.

(N3*) So, I know that I am not a BIV.

It *might* be true that one does not need to have antecedent justification to believe M3 for one's having justification to believe M1. Nevertheless, it is far from clear how one can have justification to believe N1 or N1* if one does not already presume N3 or N3*. Dogmatism of perceptual justification might be able to explain why M1's being justified does not require one to have antecedent justification of M3. But it is unclear how this view can explain the source of plausibility of the T_k employed in N1 and N1* can be gained independently of one's possession of everyday knowledge. We will probe into this point more deeply later.

The second point is, while Pryor argues that Moore's argument is not begging the question, he admits that it is *dialectically ineffective* against the sceptics in the sense that anyone who had doubts about its conclusion could not use the argument to rationally overcome those doubts. Pryor's attitude is—that is okay, as there is nothing wrong with the justificatory structure of Moore's

proof—what’s wrong is that the sceptic has doubts that he ought not to have. A dialectically ineffective argument is sufficient for solving the sceptical problem as long as it can transmit warrant from the premises to the anti-sceptical conclusion. This attitude is named ‘unambitious epistemology’ by Michael Veber, who recently provided a cogent criticism of this unambitious anti-sceptical attitude (see Veber 2019). I shall not repeat Veber’s objections here. Instead, my point is that, even Pryor cannot deny that dialectical ineffectiveness *is* indeed a deficiency of Moore’s proof (though Pryor holds that it is not a fatal one). Thus, if the gradualist anti-sceptical approach can overcome this deficiency, then this is sufficient for our purpose to show the theoretical advantage of gradualism.

One might still argue that neo-Mooreanism does not need to be formulated in the way I formulate it above; nor is neo-Mooreanism the only choice for absolutists. So, can anti-sceptics adopt Approach 1 without presuming that we have widespread knowledge? No. Remember that any anti-sceptical proposal taking Approach 1 is expected to complete two missions: 1) finding a satisfactory threshold T_k for knowledge; 2) using T_k to reject UNMET. A threshold T_k that is compatible with our lacking everyday knowledge can never be deemed ‘satisfactory’, insofar as we do not plan to surrender to scepticism. It is not only unconvincing but also inconsistent for an absolutist who chooses Approach 1 to say something like: ‘it is possible that we do not have any knowledge at all, but we will still find out a satisfactory threshold for knowledge anyway’.

Notice that I am not saying that our lack of everyday knowledge is incompatible with *the possibility that T_k exists*—after all, absolutists can argue that it is *a priori* true that knowledge entails T_k (although I am curious about how such an argument would be like), but it is still possible that we lack the widespread *empirical* knowledge that we take ourselves to have. Instead, what I am emphasising is the inconsistency between our lack of everyday knowledge and the alleged T_k ’s *being satisfactory*. The achievability of the

first mission of Approach 1 is incompatible with the possibility that we lack everyday knowledge. Precisely put, let L be 'we lack widespread everyday knowledge', and S be 'the threshold for knowledge, T_k , is satisfactory'. The statement ' $\diamond(\sim K \wedge S)$ ' is false, because a threshold is necessarily unsatisfactory if it allows the sceptical conclusion to obtain. Therefore, it is infeasible for absolutists to adopt Approach 1 without presuming the denial of the sceptical conclusion. The upshot is, any absolutist anti-sceptical proposal taking Approach 1 will either fail to offer a proper threshold for knowledge, or give rise to *petitio principia*. If it begs the question, it cannot reject scepticism appropriately; while if a proper threshold for knowledge can impossibly be given, then what is the point of choosing Approach 1? In any case, Approach 1 is not preferable. This reveals a more fundamental difficulty for any absolutist (not only neo-Mooreanists) who sticks to Approach 1.

One might object that, even if knowledge is taken to be gradable, scepticism can still argue that there are some necessary conditions for (even the most borderline cases of) knowledge that one's beliefs cannot meet. And thus the shadow of scepticism still remains within a gradualist picture. This objection is *de facto* proposing to save the sceptical argument by replacing UNMET with UNMET*. However, in section 2 we have seen that the equation between UNMET and UNMET* also obtains within a gradualist framework. Therefore, a sceptical argument based on UNMET* would fail as well if the one based on UNMET cannot succeed.

Another worry can be: Approach 2 implies gradualism, but is not gradualism also an epistemological theory about the nature of our knowledge? How can approach 2 work without presuming our possession of knowledge and thus begging the question or being dialectically ineffective? My reply is: Approach 2 can avoid *petitio principia* as it has no burden of proof for providing a satisfactory T_k , and thus it does not need to presume our possession of widespread knowledge. Gradualism blocks the sceptical argument by rejecting UNMET in a less theory-laden but more straightforward manner.

Once UNMET is rejected, then sceptics would have no reason to doubt our possession of everyday knowledge, as the standard sceptical reasoning hinges on the truth of UNMET. Moreover, unlike absolutist neo-Mooreanism, the plausibility of Approach 2 does not have to be premised on our having everyday knowledge. It is even consistent for gradualism to presume that ' $\diamond(\sim K \wedge G)$ ', where G means 'knowledge has no threshold', without preventing gradualists from rejecting UNMET (and, of course, this does not imply that gradualism concedes that scepticism is right). Therefore, Approach 2 is neither begging the question nor being subject to dialectical ineffectiveness. In summary, Approach 2 should be seen as a more ideal anti-sceptical strategy.

3. Discussions

Objection 1

Does not your argument from reference also imply that knowledge has no necessary condition? Is not this corollary too radical?

Reply 1

My argument does not entail this unacceptable corollary. Remember that the argument from reference constructs a dilemma for absolutism: if there is a threshold for knowledge, then on the one hand, it must be more lenient than SaTB/SeTB; on the other hand, it must be also *less lenient* than SaTB/SeTB. The second horn holds water when our target is T_k , because T_k ought to also serve as a sufficient condition for knowledge, but neither SaTB nor SeTB is sufficient. However, if our target is merely 'necessary condition of knowledge' in general, then the second horn would not exist, because we do not expect every necessary condition to be also sufficient for knowledge. Therefore, my argument at most implies that: 1) knowledge does not have *a certain type of* necessary condition (the type that serves as a threshold; and 2) if knowledge has a necessary condition, it should be *more lenient* than SaTB and SeTB. These two corollaries, I think, are not unacceptably radical at all.

Objection 2

But you also argued that being committed to UNMET which involves T_k is equivalent to being committed to UNMET* which involves C_k . Now that you deny the existence of T_k , is not you also denying the existence of C_k ? Does this still imply that knowledge has no necessary condition?

Reply 2

No. That is because, as we have noted before, the C_k in UNMET* is *not* a *variable* referring to any arbitrary necessary condition of knowledge. Instead, it is a very specific kind of necessary condition such that 1) sceptics can use it to deny EK, and 2) however, once it is met, then sceptics have no reason to doubt one's possession of knowledge. It serves as the last barrier that sceptics put between us and everyday knowledge. A random necessary condition of knowledge, e.g., the doxastic condition, does not fit this bill—it does not have the two aforesaid traits of C_k . For example, it is not the case that once we have beliefs, then the sceptical argument is blocked; on the other hand, sceptics can hardly deny that we have everyday knowledge by arguing that we do not have beliefs at all. Hence, rejecting knowledge's having C_k is significantly distinct from rejecting knowledge's having necessary conditions in general.

Objection 3

But how does it help to reject scepticism by simply arguing that knowledge has no threshold? 'That there is no threshold for judging whether one knows' and 'that one has widespread everyday knowledge' seem to be two irrelevant claims.

Reply 3

The focus of this chapter is the sceptical paradox. The theoretical merit of gradualism is not to entail EK from K-THRESHOLD directly. Instead, it benefits the solution to the sceptical paradox. The sceptical paradox is rooted

in the seeming contradiction between our otherwise undisputed possession of widespread everyday knowledge and the prima facie plausible sceptical argument. These are the two conflicting intuitions underpinning the sceptical paradox. The paradoxical-ness is manifested by our unwillingness to abandon either intuition given that the two intuitions are appealing. EK might be one of the most plausible claims that can function as the foundation of our epistemic conception. Hence, it is essential for the sceptical argument to be credited with adequate plausibility such that it is not too less plausible than EK to pose a paradox. In order to achieve this, sceptics need a 'scaffold' to reach the sceptical conclusion. UNMET serves as the 'scaffold' for the sceptical argument. UNMET argues that there is a T_k for knowledge, which one's belief that one is not the victim of sceptical hypotheses fails to meet. By denying the existence of T_k , gradualism rejects UNMET in an undercutting manner, as it dismantles sceptics' scaffold for reaching the sceptical conclusion. With UNMET being disproved, the sceptical argument can hardly attain sufficient plausibility to make us unwilling to abandon it easily when it contradicts EK. The sceptical paradox is thus resolved as there is no more hard choice for us.

Objection 4

Hetherington (2001; 2011) has already provided some gradualist anti-sceptical proposals. They look simpler than yours. So why is your anti-sceptical proposal still needed?

Reply 4

First of all, a small point: I do not think that my anti-sceptical proposal is obviously more complex than Hetherington's. The structures of our proposals are similar, though the differences between our arguments are salient as well. Both of our proposals target at one essential premise of the standard sceptical argument—I attempt to reject UNMET, while Hetherington mainly⁹⁶

⁹⁶ As we will see later, Hetherington (2001) also mentions the possibility of UNMET's failure. At least, if Hetherington's gradualist account of knowledge succeeds, then we have seemingly adequate reasons to believe that we *know* that we are not BIVs.

proposes to reject CLOSURE. Besides, both of our anti-sceptical proposals are accompanied by a background theory—mine appeals to K-SPECTRUM, while Hetherington's relies on minimalism (for his 2001 version of anti-scepticism) and $\leftrightarrow H$ (for his 2011 version of anti-scepticism). It is hard to say which background theory is more complicated. In any case, my anti-sceptical proposal is fundamentally distinct from his, as I reject K-THRESHOLD while Hetherington accepts it.

With these being said, now let me explain why I think Hetherington's gradualist solutions to the sceptical argument are unsuccessful. Hetherington has put forward two versions of anti-sceptical proposals along with his two versions of gradualism. Both proposals attempt to reject closure. According to the closure principle, if S knows that she is seeing zebras, then S knows that she is not dreaming of seeing zebras. Hetherington (2001) claims that this conditional is false—the antecedent's obtaining is compatible with the consequent's being false. The idea is, one's *knowing poorly* that one is seeing zebras is sufficient for satisfying the antecedent. Satisfying the consequent is just a way for one to have *much better* knowledge that one is seeing zebras. Therefore, to know (poorly) that one's seeing zebras—which is sufficient for one's knowing that one is seeing zebras—one is not required to know that one is not dreaming seeing zebras. The latter knowledge is only required when one is to know very well that one is seeing zebras.

So far so good. However, the genuine key problem here is: why is one's *knowledge that one is not dreaming* only needed if one is to know *very well* that *p*? Why is it not the case that one's knowing one is not dreaming is needed for one to know *even poorly* that one is seeing zebras? After all, what sceptics call into question is exactly whether one is able to know *even poorly* that one is seeing zebras without being able to know that one is not dreaming. Hetherington did not explicitly *prove* why. Rather, he just *assumed* that one's having justified true belief that *p* is enough for one's knowing that

p ⁹⁷. Why assume so? Hetherington (2001) provides at least two motivations—a negative one and a positive one.

The negative motivation is that sceptics have no reasonable ground for doubting this assumption. Hetherington (2001, Chap 2) diagnoses the sceptical reasoning as requiring you to know that you are not *merely* seeming not to be dreaming seeing zebras if you are to know that you are not dreaming seeing zebras. Your not *merely* seeming not to be dreaming seeing zebras includes two possibilities: 'you are not seeming not to be dreaming seeing zebras, or you are not dreaming seeing zebras' (2001: 38).

Hetherington names this disjunction as *d*. Given that the standard sceptical reasoning also accepts that you are seeming not to be dreaming, the first disjunct is ruled out; and hence sceptics are expecting you to know the second disjunct of *d*. Hetherington then argues that what sceptics actually require is: you have to know that you are not dreaming seeing zebras if you are to know that you are not dreaming seeing zebras. This requirement, according to Hetherington, is both trivial and incoherent. It is trivial because surely anything is known if it is known; it is incoherent because nothing can be known *before* it is known. In conclusion, the sceptics' requirement is unreasonable.

I do not think Hetherington gave scepticism a fair run here. First of all, one thing that we can make sure is that sceptics are not merely trivially requiring one to know the denial of sceptical hypotheses in order for one to know the denial of sceptical hypotheses. They require more substantial conditions for one to possess knowledge, for example, the *sensitivity* of one's beliefs or one's having *better evidence* to support that one is not merely dreaming. Hetherington's diagnosis neglects these independent sceptical requirements that are far more substantial than *d*. Hence, it is also not the case that

⁹⁷ Accordingly, Hetherington (2001) also argues that we *can* know that we are not dreaming seeing zebras, and thus, UNMET fails. 'I will assume that your having this justified true belief is enough for your knowing—at least in a normal, everyday way—that you are not dreaming seeing zebras' (Hetherington 2001: 37).

sceptics trivially conclude that one cannot know that one is merely dreaming just because one cannot know it *before* it is known. There are more substantial reasons that sceptics typically advocate: one's belief is insensitive; one cannot distinguish the dreaming scenarios from those not-dreaming ones; one has no better epistemic support favouring that one is not dreaming; *etc.* Therefore, it is unfair to accuse the sceptical reasoning of being trivial and incoherent—at least it is unfair to do so with the reasons that Hetherington (2001) provided.

The positive motivation, as Hetherington articulated in Chapter 3 & 4 of his 2001 book, was based on his commitment to minimalism. As we have noted before, Hetherington (2001) claims that one's true belief that p is sufficient for constituting one's minimal knowledge that p . This idea is inherited by Hetherington (2006; 2011, Ch. 2), where he interprets his minimalism in terms of his practicalism—that is, by accurately believing that p , we have already possessed the minimal knowledge that p . So scepticism can at most show that we do not know that we are not dreaming very well, but it cannot show that we do not know that we are not dreaming at all. The theoretical drawbacks of Hetherington's minimalism have been discussed in my Chapter 3. Apart from that, an additional qualm that I have for this anti-sceptical strategy is that it seems to be at odds with Hetherington's attempt to reject closure. To falsify closure, Hetherington has to prove that it is the case that the antecedent is true while the consequent is false. That is to say, Hetherington has to show that we *do not know* that we are not dreaming of seeing zebras when we know that we are seeing zebras. However, given Hetherington's minimalism, it is unclear how can we fail to know (very poorly) that we are not dreaming seeing zebras. After all, it seems that in every relevant case where we are able to know that we are seeing zebras, we would also accurately believe that we are not dreaming seeing zebras, and this should suffice to credit us with minimal knowledge that we are not dreaming seeing zebras. So it is unclear how closure's consequent can fail to be met given Hetherington's definitions of minimal knowledge.

Now we have seen that Hetherington's first version of anti-scepticism lacks proper explanations for why closure should be discarded. What about his second version of anti-sceptical proposal? In the fifth chapter of his 2011 book, Hetherington attempts to reject closure by appealing to new conceptual resources, *i.e.*, the proposed equation between knowledge that p and how-knowledge that p (which can be formulated as the biconditional ' $\leftrightarrow H$ '). Hetherington argues that, given the gradualist reading of knowledge, we should also reformulate the closure principle in a more delicate way so that it can accommodate knowledge's different degrees. Closure is reformulated as follows:

'If you know g -well that you are looking at a goldfinch, and you know perfectly well⁹⁸ that your seeing a goldfinch entails your not being a brain in a vat, then you know g -well that you are not a brain in a vat.' (2011: 196)

Anti-sceptics hold that one can know that one is looking at a goldfinch. Also, most of them would accept that, ordinarily, this knowledge is not perfect. In other words, one only knows g -well ($g < 1$; where ' $g=1$ ' means 'knows perfectly well') that one is looking at a goldfinch. However, Hetherington argues that it is not the case that one would also know g -well that one is not a BIV. Instead, he holds that we *do not know* that we are not BIVs. His argument can be summarised as follows:

- C1: If I know less than perfectly that I am seeing a goldfinch, then I know less than perfectly that I am not a BIV.
- C2: My knowledge that I am seeing a goldfinch is less than perfect, because I do not know at least one aspect of how it is that I am seeing a goldfinch—and that I am not a BIV can be this unknown aspect.
- C3: Therefore, C1 is false. My knowing I am seeing a goldfinch is compatible with my not knowing that I am not a BIV.

⁹⁸ Hetherington (2011: fn 45) admits that he uses 'perfectly well' here just in order to prevent the formulation from being too complicated.

C1 is an application of the reformulated closure (assuming that I know perfectly well that my seeing a goldfinch entails that I am not a BIV). C2 explains why my everyday knowledge is imperfect—it might be because I do not know that I am not a BIV. Therefore, C3 follows in that it can be the case that the antecedent of C1 is true while its consequent is false. Hetherington then proposes to generalise this reasoning such that it can be applied to disarm more forms of scepticism that do not employ the closure principle. He constructs a dilemma for sceptics:

For horn 1, which assumes that \sim BIV is a part of a minimal truthmaker for 'I am seeing a goldfinch', Hetherington argues that 'there can be parts of a minimal truthmaker for p that are not known, even as p itself is known — hence, even as a minimal truthmaker for p is known' (2011: 197). This explains why one's knowledge that p is imperfect. Therefore, \sim BIV can be such an unknown part of a minimal truthmaker for 'I am seeing a goldfinch'. It is only needed when one is to know perfectly well that one is seeing a goldfinch. Knowing \sim BIV can indeed *improve* my knowledge that I am seeing a goldfinch, while failing to have the former knowledge does not deprive me of my latter knowledge.

For horn 2, which assumes that \sim BIV is *not* a part of a minimal truthmaker for 'I'm seeing a goldfinch', Hetherington asserts that '[k]nowledge that p need not include knowledge of anything that is not part of a minimal truthmaker for p ' (2011: 200). Therefore, one's knowledge that \sim BIV is still not required for one's knowing that 'I'm seeing a goldfinch'.

There are three reasons why I find Hetherington's second gradualist solution to the sceptical problem is also unsuccessful. First, it is obvious that this proposal hinges essentially on his analysis of knowledge that p in terms of how-knowledge that p . As we have seen in Chapter 3, his argument for ' \leftrightarrow H'

is problematic. Consequently, the corresponding anti-sceptical proposal is ill-grounded.

Second, I am afraid that the argument against closure is dialectically ineffective. C2 is logically equivalent to 'I know less than perfectly that I am seeing a goldfinch, and I do not know that I am not a BIV'—which is simply the negation of C1, not to mention that the first conjunct just assumes the failure of the sceptical conclusion. The key explanandum here is why I can know a minimal truthmaker for 'I am seeing a goldfinch' without knowing that 'I am not a BIV', if the latter is a part of the former. Hetherington seems to imply that there is at least one minimal truthmaker⁹⁹ for 'I am seeing a goldfinch' that does not involve \sim BIV as an indispensable part. But he does not articulate what this minimal truthmaker is, or why such a minimal truthmaker exists. To explain this point, what Hetherington (2011: 198) uses is an analogy that your knowledge of there being something physical in front of you is not required as part of knowing that there is a frog there. However, notice that champions of closure do not need to hold that the former knowledge *is a part of* the latter knowledge. They do not need to presume any inclusion relation between the two items of knowledge. They just need to show that if one possesses the latter knowledge, then one should also possess the former one. Accordingly, proponents of competent deduction closure (see footnote 85, p. 213) would object that it is unclear how one can know that 'there is a frog in front of me' without thereby knowing that 'there is something physical in front of me' if one can competently deduce the latter from the former.

Last but not least, similar to his first version of gradualist anti-sceptical proposal, Hetherington's second proposal also gives rise to a potential incoherence. Given that one's knowing a minimal truthmaker for p is sufficient for one to know that p , it seems that we *do know* that we are not BIVs. That is

⁹⁹ Hetherington suggests that a proposition can have more than one minimal truthmaker, and one only needs to know a minimal truthmaker for p in order to possess the minimal knowledge that p .

because if I know that I am seeing a goldfinch and that my seeing a goldfinch entails that I am not a BIV (as the antecedent of closure requires), then at least I know that 'I am seeing a goldfinch' obtains, and that 'if I am seeing a goldfinch, then I am not a BIV' obtains—remember that Hetherington suggests that the minimal aspect of how it is p is 'that p obtains' (see Hetherington 2011: 174). 'If I am seeing a goldfinch, then I am not a BIV' is identical with a disjunction 'I am not seeing a goldfinch, or I am not a BIV'. Now that I know that the negation of the first disjunct obtains, I should also know that the second disjunct obtains. Therefore, I should know a minimal truthmaker for 'I am not a BIV', and thus, according to Hetherington's minimalism, know that I am not a BIV. If that is the case, then C2 is incompatible with its background theory, to wit, Hetherington's gradualist account of knowledge that p in terms of how-knowledge that p .

4. Concluding Remarks

The debate between epistemic absolutism and gradualism is meaningful, and research on that debate can be fruitful. This chapter illustrates how gradualism can prevail over absolutism in terms of its application in solving the sceptical paradox. Within the orthodox absolutist picture, the sceptical paradox is premised on an intuitively appealing claim that one's belief that 'I am not a BIV' does not meet the threshold for knowledge. An absolutist anti-sceptical approach rejects this claim by proving that there is a threshold that is met, while a gradualist approach straightforwardly denies that knowledge has a threshold. It has been shown that the absolutist approach succumbs to the problem of reference and the problem of *petitio principia* and thus ought to be discarded. In contrast, by interpreting knowledge as a spectrum concept, gradualism can block the sceptical argument in a more undercutting way without giving rise to those problems that absolutism is subject to. In light of this, gradualism enjoys advantages over absolutism when it comes to solving the sceptical problem.

Conclusion: Towards the Concept of Gradable Knowledge

A key mission of epistemology is to research into the properties of knowledge, in particular, propositional knowledge. The gradability is an important property that worthy of discussions. The standard view, namely, *epistemic absolutism*, holds that to know a proposition is an absolute yes-or-no affair. In other words, propositional knowledge is ungradable. The main reason for accepting absolutism is that we usually do not use 'knows' as a gradable predicate in the ordinary language. In addition, some advocates of absolutism treat propositional knowledge as an ungradable concept because its putative object, *viz*, the truth of a proposition, is ungradable. Besides, some people refuse to view knowledge as gradable because they hold that this would perforce lead to contextualism, which is an epistemological theory that they find independently implausible.

This thesis starts with a thorough objection to this long-held absolutist view. It has been revealed that epistemic absolutism is not as well supported by arguments as orthodoxy has it:

Firstly, the linguistic argument for absolutism has been debunked by both the internal criticism and the external criticism. The internal criticism shows that linguistic data against epistemic absolutism can also be found in both English and non-English languages, and thus it is unfair to assert that linguistic evidence favours absolutism. The external criticism calls into question the linguistic methodology employed by absolutists. It has been shown that if our linguistic usages of 'knows that' in the ordinary language determine the conceptual nature of knowledge *per se*, then knowledge should also be non-factive given that we sometimes use knowledge in a non-factive manner. An analysis of a non-doxical account of knowledge and hinge epistemology also illustrates that there is a gap between the linguistic usages of 'knows' and the conceptual essence of knowledge.

Secondly, the argument from object has been proven to be untenable as it is premised on the problematic KT-schema, *i.e.*, to know that p is just to know that p is true. By revealing a paradox that KT-schema would give rise to and providing counterexamples to KT-schema, I have demonstrated that the object of knowledge-that- p is more than the truth of p . It has to at least include an adequate degree of understanding of p 's semantic meaning. Apart from this, even if one granted KT-schema, it would be a *non sequitur* to entail the ungradability of an epistemic concept from the ungradability of its object. Otherwise, it would follow that some paradigmatic gradable concepts such as 'belief' and 'remember' are also ungradable.

Finally, it has been argued that we do not need to endorse contextualism to see knowledge as gradable. Hence the opponents of contextualism do not need to be committed to absolutism.

So, if epistemic absolutism is ill-grounded, what kind of view shall we embrace instead? The answer given by this thesis is *epistemic gradualism*, a view that propositional knowledge is gradable. The most prominent advocate of this view is Stephen Hetherington. He developed two gradualist proposals, the first one is based on his anti-justificationism (a view that justification is not necessary for knowledge), the second one hinges on his how-knowledge theory. According to Hetherington, the central divergence between gradualism and absolutism is whether one can be recognised as knowing that p better by virtue of better justification. He holds that absolutists would refuse to admit that one's knowledge can be improved by better epistemic support, and hence absolutism is obviously implausible. But he grants that there is a clear cut-off point distinguishing knowledge from everything that falls short of knowledge. In other words, he grants that there is a threshold dividing knowledge and ignorance. In summary, Hetherington is advocating a version of local gradualism. That is, he supports beyond-threshold gradualism, but rejects threshold gradualism (see Chapter One).

It has been shown that Hetherington's characterisation of the gradualism/absolutism debate is problematic. First, it misdiagnoses the key divergence between gradualism and absolutism and thus gives rise to an equivocal attitude towards knowledge's gradability. In fact, many absolutists would be glad to endorse beyond-threshold gradualism, to wit, knowledge can be seen as gradable in the sense that it can be improved by virtue of better justification. The problem lies in that people do not therefore treat knowledge as an overall gradable concept. The idea is, for the equivocal attitude, knowledge is gradable in the beyond-threshold sense, but it is ungradable in a more general sense. The second defect of the present model of the debate between gradualism and absolutism is that Hetherington's gradualist proposals are not appealing enough to constitute a balanced debate. The anti-justificationist ground of his first gradualist proposal invites an unnecessarily heavy burden of proof for gradualism, and his second proposal is based on an invalid argument equating knowledge that p to how-knowledge that p .

This thesis has remodelled the debate between gradualism and absolutism. It has been argued that the central divergence of the two views should be summarised as the competition between two readings of knowledge. They are:

[K-THRESHOLD]

Knowledge is a threshold concept in the sense that there is a threshold T_k , which distinguishes knowledge from everything that falls short of knowledge.

[K-SPECTRUM]

Knowledge is a spectrum concept in the sense that there is *not* a threshold T_k distinguishing knowledge from everything that falls short of knowledge, but only better or

worse instances of knowledge that can be graded into different degrees on the relevant scale S.

By endorsing K-SPECTRUM, one rejects both beyond-threshold absolutism and, more importantly, threshold absolutism (which can be further subdivided into trans-threshold absolutism and intra-threshold absolutism, see Chapter Three). The reconstructed model of the debate depicts the genuine divergence between gradualism and absolutism more accurately and more explicitly. In addition, it can avoid the defects of Hetherington's gradualism and thereby yield a more balanced and attractive debate.

Being based on K-SPECTRUM, the reconstructed framework of gradualism is also of greater flexibility. The gradability of knowledge can be interpreted from various perspectives in terms of different relevant scales. By way of illustration, drawing on Sosa's virtue epistemology, one can account for knowledge's gradability in terms of the degrees of accuracy (for example, by employing a fuzzy logic), adroitness (to what degree one's true belief manifests one's cognitive competence), and adptness (to what degree one's belief's being true is attributable to the manifestation of one's relevant cognitive competence) of one's belief. The upshot is, gradualism does not need to be bound by any specific account of knowledge. Instead, any full account of knowledge would need to register how propositional knowledge is a matter of degree.

One significant goal of this thesis is to overcome the equivocal attitude towards knowledge's gradability. The strategy that I took consists of two steps. First, to undermine absolutism, as noted before, this thesis scrutinised the arguments for absolutism from three perspectives. It has been shown that all three mainstream arguments for absolutism are untenable. Second, to motivate gradualism, it has been demonstrated that gradualism would enable us to solve many epistemological issues that absolutism cannot properly address. I have demonstrated several theoretical advantages that gradualism

enjoys over absolutism. These include gradualism's application in providing a unified account of different varieties of knowledge, which eliminates the putative asymmetry problem (see Chapter Four). Besides, it has been shown that gradualism performs better when it comes to the reconciliation of the two conflicting intuitions about knowledge's compatibility of environmental luck, and accounting for the puzzling phenomena of faultless disagreements (see Chapter Five). In closing, the last chapter has explored epistemic gradualism's application in resolving the sceptical paradox. It has been concluded that knowledge is best to be viewed as a spectrum concept, which enables a more satisfactory reply to scepticism. Therefore, gradualism ought to be seen as preferable to absolutism, and there would be no reason to stick to the absolutist orthodoxy.

Albeit currently underexplored, the debate between absolutism and gradualism is epistemologically valuable. However, the extant model of the debate interprets the core divergence of gradualism and absolutism in a problematic way, and thus might give rise to the equivocal attitude. This thesis suggests that we should remodel the debate by focusing on the competition between the threshold reading and the spectrum reading of knowledge. This can help us to construct a more philosophically meaningful and fruitful debate. I believe that the application of gradualism is not limited to aforementioned issues that have been discussed in this thesis. Some other relevant questions worth exploring include: How can gradualism be integrated with the extended knowledge theory? Would the genealogy of knowledge favour gradualism over absolutism? If we reject the idea of minimal knowledge, is there maximal knowledge? Which extant account of knowledge best tallies with gradualism?

Interesting as these questions are, this thesis can only leave them open for now due to the limit of space. Also, I am far from believing that this thesis has made a conclusive case for defending gradualism and rejecting absolutism. Instead, I prefer to see this thesis as reopening rather than concluding the

gradualism/absolutism debate. Further discussion on this debate will have a bright and fruitful prospect that is worthy of deeper exploration.

Bibliography

- Adams, F. R. & Aizawa, K. (2010). Defending the bounds of cognition. In Richard Menary (ed.), *The Extended Mind*. pp. 67-80. MIT Press.
- Armstrong, D. M. (2004). *Truth and Truthmakers*. Cambridge University Press.
- Ashton, N. A. (2015). Undercutting Underdetermination-Based Scepticism. *Theoria* 81 (4): 333-354.
- Ayer, A. J. (1935). The Criterion of Truth. *Analysis* 3: 28-32.
- Bac, M. & Irmak, N. (2011). Knowing Wrongly: An Obvious Oxymoron, or a Threat for the Alleged Universality of Epistemological Analyses? *Croatian Journal of Philosophy* 11 (3): 305-321.
- Båve, A. (2013). Formulating deflationism. *Synthese* 190 (15):3287-3305.
- Beck, S. & Yael S. (2002). Pluralities of Questions. *Journal of Semantics* 19, no. 2: 105–57.
- Berman, S. 1987. Situation-Based Semantics for Adverbs of Quantification. *University of Massachusetts Occasional Papers* 12: 45-68.
- Bernecker, S. (2012). Sensitivity, safety, and closure. *Acta Analytica*, 27(4): 367-381.
- Beall, J. C. (2001). Is Yablo's paradox non-circular?. *Analysis*, 61: 176-87.
- Bengson, J. & Moffett, M. A. (2011). Nonpropositional Intellectualism. In John Bengson & Marc A. Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*. pp. 161-195. Oxford University Press.
- Bierwisch, M. (1989). The semantics of gradation. In Manfred Bierwisch & Ewald Lang (eds.), *Dimensional adjectives: Grammatical structure and conceptual interpretation*. pp. 71– 261. Berlin: Springer-Verlag.
- BonJour, L. (2002). *Epistemology: Classical Problems and Contemporary Responses*. Lanham MD: Rowman and Littlefield.
- BonJour, L. (2003). A Version of Internalist Foundationalism. In BonJour and Ernest Sosa (eds.), *Epistemic Justification*. pp. 21-3. Maiden, Mass.: Blackwell.
- BonJour, L. (2010). The Myth of knowledge. *Philosophical Perspectives* 24 (1): 57-83.

- Bogardus, T. (2012). What Certainty Teaches. *Philosophical Psychology* 25 (2): 227 - 243.
- Bogardus, T. & Marxen, C. (2014). Yes, Safety is in Danger. *Philosophia* (2): 1-14.
- Boghossian, P. (2006). *Fear of Knowledge: Against Relativism and Constructivism*. Oxford University Press.
- Bovell, C. R. (2012), *Ideas at the Intersection of Mathematics, Philosophy, and Theology*, Eugene: Wipf and Stock
- Broncano-Berrocal, F. (2014). Is Safety In Danger? *Philosophia* 42 (1): 1-19.
- Brower, B. W. (2004). Good Knowledge, Bad Knowledge. *International Philosophical Quarterly* 44 (1): 107-108.
- Buckwalter, W. (2014). Factive Verbs and Protagonist Projection. *Episteme* 11 (4): 391-409.
- Byrne, A. (2004). How Hard Are the Sceptical Paradoxes? *Noûs* 38 (2): 299-325.
- Cappelen, H. & Hawthorne, J. (2009). *Relativism and Monadic Truth*. Oxford University Press UK.
- Carr, D. (1981). Knowledge in Practice. *American Philosophical Quarterly* 18 (1):53 - 61.
- Carter, J. A. (2013). A Problem for Pritchard's Anti-luck Virtue Epistemology. *Erkenntnis* 78 (2): 253-275.
- Carter, J. A. & Pritchard, D. (2015). Knowledge-How and Epistemic Luck. *Noûs* 49 (3): 440-453.
- Carter, J. A. (2016). *Metaepistemology and Relativism*. Basingstoke: Palgrave Macmillan.
- Carter, J. A.; Pritchard, D. & Shepherd, J. (2019). Knowledge-how, understanding-why, and epistemic luck: an experimental study. *Review of Philosophy and Psychology*: 1-34.
- Cath, Y. (2011). Knowing How Without Knowing That. In John Bengson & Mark Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*. Oxford University Press.

- Chisholm, R. M. (1966). *Theory of Knowledge*. Englewood Cliffs, N.J., Prentice-Hall.
- Church, Ian M. (2012). Epistemic Contextualism, Epistemic Relativism, and Disagreement: Reply to Robin McKenna. *Philosophical Writings*: 100-103.
- Clark, A. & Chalmers, D. J. (1998). The Extended Mind. *Analysis* 58 (1):7-19.
- Clark, M. (2007). *Paradoxes from A to Z*, Routledge: Oxford and New York
- Clarke, R. (2013). Belief Is Credence One (in Context). *Philosophers' Imprint* 13:1-18.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives* 13 (s13): 57-89.
- Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research* 65: 309–29.
- Colao, D., Buckwalter, W., & Stich, S. (2012). Epistemic Intuitions in Fake-Barn Thought Experiments (September 23, 2012). Available at SSRN: <http://ssrn.com/abstract=1973351>
- Crane, T. (2012). Tye on Acquaintance and the Problem of Consciousness. *Philosophy and Phenomenological Research* 84 (1): 190-198.
- Curry, H. (1942). The Inconsistency of Certain Formal Logics, *Journal of Symbolic Logic* 7: 115-117.
- Davies, A. (2018). Infallibilism and Easy Counter-Examples, *Grazer Philosophische Studien* 95 (4):475-499.
- Davis, W. A. (2007). Knowledge claims and context: loose use. *Philosophical Studies* 132 (3): 395-438.
- DePaul, M. R. (2009). Ugly Analyses and Value. In Adrian Haddock, Alan Millar & Duncan Pritchard (eds.), *Epistemic Value*. pp. 112-138. Oxford University Press.
- DeRose, K. (1995). Solving the Skeptical Problem. *Philosophical Review* 104 (1): 1-52.
- DeRose, K. (2004). Single Scoreboard Semantics. *Philosophical Studies* 119 (1-2): 1-21.
- Dodd, D. (2011). Against Fallibilism, *Australasian Journal of Philosophy* 89 (4): 665-685.

- Domaneschi, F. & Di Paola, S. (2019). Relevance and Non-Factive Knowledge Attributions. *Acta Analytica* 34 (1): 83-115.
- Dummett, M. (1959). Truth. *Proceedings of the Aristotelian Society*, n.s. 59: 141-162.
- Dutant, J. (2007). The Limited Gradability of 'Know', Unpublished manuscript. Available at http://julien.dutant.free.fr/papers/JDutant_TheLimitedGradabilityOfKnow.pdf
- Dutant, J. (2016). How to Be an Infallibilist, *Philosophical Issues* 26 (1): 148-171.
- Dretske, F. (1970). Epistemic Operators. *The Journal of Philosophy* 67: 1007-1023.
- Dretske, F. (1971). Conclusive Reasons, *Australasian Journal of Philosophy* 49: 1-22
- Dretske, F. (1981). The Pragmatic Dimension of Knowledge. *Philosophical Studies* 40 (3): 363-378.
- Dretske, F. (2014). The Case against Closure. In M. Steup & Ernest Sosa (eds.), *Contemporary Debates in Epistemology*. pp. 13-25. Malden, Ma: Blackwell.
- Engel Jr, M. (2004). What's Wrong with Contextualism, and a Noncontextualist Resolution of the Skeptical Paradox. *Erkenntnis* 61 (2-3): 203-231.
- Fantl, J. & McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford University Press.
- Farkas, K. (2015). Belief May Not Be a Necessary Condition for Knowledge. *Erkenntnis* 80 (1): 185-200.
- Field, H. (1994). Disquotational Truth and Factually Defective Discourse, *Philosophical Review* 103 (3): 405-52.
- Field, H. (2008). *Saving Truth from Paradox*, Oxford: Oxford University Press.
- Feldman, R. (2002). Review of Stephen Hetherington, *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*. *Notre Dame Philosophical Reviews* 2002 (8).

- Foley, R. (2009). Beliefs, Degrees of Belief, and the Lockean Thesis. In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*. pp. 37-47. Springer.
- Gendler, T. S. & Hawthorne, J. (2005). The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies. *Philosophical Studies* 124 (3): 331-352.
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis* 23 (6): 121-123.
- Goldberg, S. (2015). Epistemic Entitlement and Luck. *Philosophy and Phenomenological Research* 91 (2): 273-302.
- Goldman, A. I. (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy* 73 (November): 771-791.
- Glock, H.-J. (1996). *A Wittgenstein Dictionary*. Oxford: Blackwell.
- Goldman, A. I. (1967). A Causal Theory of Knowing. *The Journal of Philosophy*, 64(12): 357-372.
- Goldman, A. I. (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy* 73 (November): 771-791.
- Goldman, A. I. (1999). *Knowledge in a Social World*. Oxford University Press.
- Greco, D. (2015). How I Learned to Stop Worrying and Love Probability 1. *Philosophical Perspectives* 29 (1): 179-201.
- Greco, J. (2009). 'The Value Problem', In Adrian Haddock, Alan Millar & Duncan Pritchard (eds.), *Epistemic Value*. pp. 313-321. Oxford University Press.
- Greco, J. (2010). *Achieving Knowledge*, Cambridge: Cambridge University Press
- Groenendijk, J. & Stokhof, M. (1982). Semantic Analysis of Wh-Complements. *Linguistics and Philosophy* 5 (2): 175 - 233.
- Grundmann, T. (2018). Saving Safety From Counterexamples. *Synthese*.
- Hannon, M. (2013). 'Knows' Entails Truth. *Journal of Philosophical Research* 38: 349-366.
- Hannon, M. (2017). A Solution to Knowledge's Threshold Problem. *Philosophical Studies* 174 (3): 607-629.

Harrison, J. (1978). IX—If I Know, I Cannot Be Wrong. *Proceedings of the Aristotelian Society* 79 (1):137-150.

Hawthorne, J. (2002). Deeply Contingent A Priori Knowledge, *Philosophy and Phenomenological Research*, 65(2): 247-69.

Hawthorne, J. (2003). *Knowledge and Lotteries*. Oxford University Press.

Hazlett, A. (2010). The Myth of Factive Verbs. *Philosophy and Phenomenological Research* 80 (3): 497-522.

Hazlett, A. (2012). Factive Presupposition and the Truth Condition on Knowledge. *Acta Analytica* 27 (4): 461-478.

Hazlett, A. (2014). *A Critical Introduction to Skepticism*. Bloomsbury Academic.

Hetherington, S. (1999). Knowing Failably. *Journal of Philosophy* 96 (11): 565-587.

Hetherington, S. (2001). *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*. Oxford University Press.

Hetherington, S. (2005). Knowing (How It Is) That P: Degrees and Qualities of Knowledge. *Veritas—Revista de Filosofia da Pucrs* 50 (4): 129-152.

Hetherington, S. (2006). Knowledge's Boundary Problem. *Synthese* 150 (1): 41-56.

Hetherington, S. (2007). Is This a World Where Knowledge Has to Include Justification? *Philosophy and Phenomenological Research* 75 (1): 41-69.

Hetherington, S. (2011). *How to Know: A Practicalist Conception of Knowledge*. Wiley-Blackwell.

Hetherington, S. (2013). Knowledge Can Be Lucky. In Matthias Steup & John Turri (eds.), *Contemporary Debates in Epistemology*. 164-176. Blackwell.

Hiller, A. & Neta, R. (2007). Safety and Epistemic Luck. *Synthese* 158 (3): 303-313.

Holton, R. (1997). Some Telling Examples: A Reply to Tsohatzidis. *Journal of Pragmatics* 28: 625-628.

Horwich, P. (1998). *Truth*. Oxford: Blackwell (first edition 1990).

Horwich, P. (2000). The Sharpness of Vague Terms. *Philosophical Topics*, 28(1): 83-92.

- Horwich, P. (2010). *Truth-Meaning-Reality*. Oxford: Clarendon Press.
- Hu, I. (2015). Epistemicism, paradox, and conditional obligation. *Philosophical Studies* 172 (8): 2123-2139.
- Hudson, R. (2014). Saving Pritchard's anti-luck virtue epistemology: the case of Temp. *Synthese* 191 (5): 1-15.
- Hyde, D. (2010). Vagueness, Logic and Ontology, *Bulletin of Symbolic Logic* 16 (4):531-533.
- Immerman, D. (2016). Sensitivity, Reflective Knowledge, and Skepticism. *International Journal for the Study of Skepticism* 6 (4): 351-367.
- Jaster, R. (2013). Contextualism and Gradability. *GAP.8 Proceedings*: 319-323.
- Jespersen, B. (2008). Knowing that p rather than q . *Sorites* 20: 125-134.
- Jorgenson, J. (1953). Some Reflections on Reflexivity. *Mind*, Vol. 62, No. 247: 289-300.
- Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge University Press.
- Kelp, C. (2009). Knowledge and Safety. *Journal of Philosophical Research* 34: 21-31.
- Kelp, C. (2013). Knowledge: The Safe-apt View. *Australasian Journal of Philosophy*, 91: 265-278.
- Kelp, C. (2017). Knowledge First Virtue Epistemology. In Adam Carter, Emma Gordon & Benjamin Jarvis (eds.), *Knowledge First: Approaches in Epistemology and Mind*. Oxford University Press.
- Kelp, C. (2018). *Good Thinking. A Knowledge First Virtue Epistemology*. London, UK: Routledge.
- Klein, P. (2004). Closure Matters: Academic Skepticism and Easy Knowledge. *Philosophical Issues*, 14: 165-184.
- Kölbel, M. (2004). Faultless Disagreement. *Proceedings of the Aristotelian Society* 104 (1): 53-73.
- Kompa, N. (2015). Contextualism and Disagreement. *Erkenntnis* 80 (1): 137-152.

- Kripke, S.A. (1975). Outline of a Theory of Truth, *Journal of Philosophy*, 72(19): 690-716.
- Kripke, S. A. (2011). Nozick on Knowledge, in *Philosophical Troubles: Collected Papers, Volume 1*, pp. 162-224. Oxford: Oxford University Press,
- Kusch, M. (2011). Knowledge and Certainties in the Epistemic State of Nature. *Episteme* 8 (1): 6-23.
- Kusch, M. & McKenna, R. (2018). The Genealogical Method in Epistemology. *Synthese*.
- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Lackey, J. (2006). Pritchard's Epistemic Luck. *Philosophical Quarterly* 56 (223): 284-289.
- Lahiri, U. (2000). Lexical Selection and Quantificational Variability in Embedded Interrogatives. *Linguistic and Philosophy* 23 (4): 325–489.
- Lai, C. (2019). Is 'Knowing that P' Identical with 'Knowing that "P" Is True'? *Philosophia*: 1-18.
- Le Morvan, P. (2002). Is Mere True Belief Knowledge? *Erkenntnis* 56 (2): 151-168.
- Le Morvan, P. (2005). Goldman on Knowledge as True Belief. *Erkenntnis* 62 (2): 145-155.
- Lehrer, K. (2000). Sensitivity, Indiscernibility and Knowledge. *Philosophical Issues* 10: 33-37.
- Leite, A. (2004). Skepticism, Sensitivity, and Closure, or Why the Closure Principle Is Irrelevant to External World Skepticism. *Croatian Journal of Philosophy* 12 (3): 335-350.
- Lewis, David (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic* 8 (1): 339-359.
- Littlejohn, C. (2008). From E = K to Scepticism?, *Philosophical Quarterly* 58 (233): 679-684.
- López de Sa, D. (2008). Presuppositions of Commonality: An Indexical Relativist Account of Disagreement. In G. García-Carpintero & M. Koelbel (eds.), *Relative Truth*. pp. 297-310. Oxford University Press.

- Lycan, W. G. (2006). On the Gettier Problem Problem. In Stephen Hetherington (ed.), *Epistemology Futures*. pp. 148-168. Oxford University Press.
- MacFarlane, J. (2007). Relativism and Disagreement. *Philosophical Studies* 132 (1): 17-31.
- MacFarlane, J. (2009). Nonindexical Contextualism. *Synthese* 166 (2): 231-250.
- Matthews, J. (2013). Ungrounded Self-reference: the Solution to the Liar Paradox?
- McGinn, M. (1989). *Sense and Certainty: A Dissolution of Scepticism*. Oxford: Blackwell.
- McGinn, M. (2010). Ludwig Wittgenstein. In R. Neta & D. H. Pritchard (Eds.), *The Routledge companion to epistemology*, pp. 763-773, London: Routledge.
- McKenna, R. (2012). Epistemic Contextualism, Epistemic Relativism and Disagreement. *Philosophical Writings*.
- Moffett, M. A. (2003). Knowing Facts and Believing Propositions: A solution to the problem of doxastic shift. *Philosophical Studies* 115 (1): 81-97.
- Moser, P. K. (1987). Propositional Knowledge. *Philosophical Studies* 52 (1): 91-114.
- Moyal-Sharrock, D. (2004). *Understanding Wittgenstein's on certainty*. London: Palgrave Macmillan
- Myers-Schulz, B. & Schwitzgebel, E. (2013). Knowing That P without Believing That P. *Noûs* 47 (2): 371-384.
- Neta, R. (2002). S Knows that P. *Noûs* 36 (4): 663-681.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford University Press.
- Pavese, C. (2017). Know-How and Gradability. *Philosophical Review* 126 (3): 345-383.
- Plunkett, D. & Sundell, T. (2013). Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint* 13.
- Priest, G. (1984). Logic of Paradox Revisited. *Journal of Philosophical Logic* 13 (2): 153-179.
- Priest, G. (1997). Yablo's Paradox. *Analysis* 57: 236-42.

- Priest, G. (2006). *In Contradiction (2nd ed.)*, Oxford: Oxford University Press.
- Priest, G. (2008). *An Introduction to Non-Classical Logic*, Cambridge: Cambridge University Press, second edition.
- Prior, A. N., (1955). Curry's Paradox and 3-Valued Logic, *Australasian Journal of Philosophy* 33: 177-82.
- Pritchard, D. (2005). The Structure of Sceptical Arguments. *Philosophical Quarterly* 55 (218): 37-52.
- Pritchard, D. (2007). Anti-Luck Epistemology. *Synthese* 158: 277-297.
- Pritchard, D. (2008). Sensitivity, Safety, and Anti-luck Epistemology. In John Greco (ed.), *The Oxford Handbook of Scepticism*. pp. 437-455. Oxford: Oxford University Press.
- Pritchard, D. (2009a). The Value of Knowledge. *The Harvard Review of Philosophy* 16 (1): 86-103.
- Pritchard, D. (2009b). Safety-Based Epistemology: Whither Now? *Journal of Philosophical Research* 34: 33-45.
- Pritchard, D. (2010). The Value of Knowledge: Understanding. In Haddock, Adrian, Alan Millar and Duncan Pritchard (ed.). *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press.
- Pritchard, D. (2011). Wittgenstein on Scepticism. In Marie McGinn & Oskari Kuusela (eds.), *The Oxford Handbook of Wittgenstein*. pp. 523-549. Oxford University Press.
- Pritchard, D. (2012a). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109 (3): 247-279.
- Pritchard, D. (2012b). Wittgenstein and the Groundlessness of Our Believing. *Synthese* 189 (2): 255-272.
- Pritchard, D. (2015a). Anti-luck Epistemology and the Gettier Problem. *Philosophical Studies* 172 (1): 93-111.
- Pritchard, D. (2015b). *Epistemic Angst: Radical Scepticism and the Groundlessness of Our Believing*. Princeton: Princeton University Press.
- Pritchard, D., Alan M., & Adrian H., (2010). *The Nature and Value of Knowledge: Three Investigations*. Oxford, UK: Oxford University Press.

- Pryor, J. (2004). What's Wrong with Moore's Argument? *Philosophical Issues* 14 (1):349–378.
- Pryor, J. (2012). When Warrant Transmits. In A. Coliva (ed.), *Wittgenstein, Epistemology and Mind: Themes from the Philosophy of Crispin Wright*, pp. 269–303. Oxford: Oxford University Press.
- Quine, W.V.O. (1970). *Philosophy of Logic*, Englewood Cliffs: Prentice Hall.
- Ramsey, F. P. (1927). Facts and Propositions, *Proceedings of the Aristotelian Society* 7 (Supplementary): 153-170.
- Recanati, F. (2008). Moderate relativism. In M. García-Carpintero & M. Kölbel (Eds.), *Relative truth*, pp. 41–62. Oxford: Oxford University Press.
- Roland, J. (1958). On 'Knowing How' and 'Knowing That'. *Philosophical Review* 67 (3): 379-388.
- Rothschild, D. & Spectre, L. (2018). At the Threshold of Knowledge. *Philosophical Studies* 175 (2): 449-460.
- Ryan, S. (2017). Why Knowledge is Special. *Philosophy* 92 (2): 249-269.
- Ryle, G. (1949). *The Concept of Mind*. Hutchinson & Co.
- Sartwell, C. (1991). Knowledge is Merely True Belief. *American Philosophical Quarterly* 28 (2): 157-165.
- Sartwell, C. (1992). Why Knowledge Is Merely True Belief', *The Journal of Philosophy* 89(4): 167-180.
- Schönbaumsfeld, G. (2017). The 'Default View' of Perceptual Reasons and 'Closure-Based' Sceptical Arguments. *International Journal for the Study of Skepticism* 7 (2): 114-135.
- Seidel, M. (2014). *Epistemic Relativism. A Constructive Critique*. Palgrave-Macmillan.
- Sgaravatti, D. & Zardini, E. (2008). Knowing How to Establish Intellectualism. *Grazer Philosophische Studien* 77 (1): 217-261.
- Shapiro, S. (2006). *Vagueness in Context*, Oxford University Press
- Shope, R. K. (1983). The Analysis of Knowing: A Decade of Research. *Philosophy and Phenomenological Research* 46 (3): 523-528.
- Simion, M. & Kelp, C. (2016). The Tertiary Value Problem and the Superiority of Knowledge. *American Philosophical Quarterly* 53 (4):397-410.

- Skidmore, A. (1993). Is Knowledge Merely True Belief?, *Southwest Philosophy Review* 9(1), 71-76.
- Smith, N.J.J. (2008). *Vagueness and Degrees of Truth*. Oxford University Press.
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: University Press UK.
- Sorensen, R. (1998). Yablo's Paradox and Kindred Infinite Liars. *Mind* 107: 137-155.
- Sosa, E. (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives* 13: 141-154.
- Sosa, E. (2001). Human Knowledge, Animal and Reflective. *Philosophical Studies* 106 (3): 193-196.
- Sosa, E. (2004). Replies. In John Greco (ed.), *Ernest Sosa: And His Critics*, pp. 275-327. Malden MA: Blackwell.
- Sosa, E. (2007). *A Virtue Epistemology*. Oxford University Press.
- Sosa, E. (2009). *Reflective Knowledge*. Oxford University Press.
- Sosa, E. (2015). *Judgment & Agency*. Oxford University Press UK.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2011). *Know How*. Oxford University Press.
- Stroud, B. (1984). *The Significance of Philosophical Scepticism*. Oxford University Press.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2011). *Know How*. Oxford University Press.
- Tarski, A. (1933). The Concept of Truth in the Languages of the Deductive Sciences (Polish), *Prace Towarzystwa Naukowego Warszawskiego, Wydział III Nauk Matematyczno-Fizycznych*, 34, Warsaw; reprinted in Zygmunt 1995: 13--172; expanded English translation in Tarski 1983: 152--278.
- Tarski, A. (1944). The Semantic Conception of Truth, *Philosophy and Phenomenological Research*, IV: 341-75.

- Tarski, A. (1983). *Logic, Semantics, Metamathematics: Papers from 1923 to 1938*, edited with introduction and analytic index by John Corcoran, Indianapolis: Hackett Publishing Company.
- Tsohatzidis, S. L. (2012). How to Forget that 'Know' is Factive. *Acta Analytica* 27 (4):449-459.
- Turri, J. (2011a). Mythology of the Factive. *Logos & Episteme* 2(1): 141-150.
- Turri, J. (2011b). Manifest Failure: The Gettier Problem Solved. *Philosophers' Imprint* 11.
- Unger, P. (1971). A Defense of Skepticism. *Philosophical Review* 80 (2): 198-219.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford University Press.
- Vogel, J. (1987). Tracking, Closure and Inductive Knowledge. In S. Luper-Foy (ed.), *The Possibility of Knowledge: Nozick and His Critics*. pp. 197-215. Totowa: Rowman and Littlefield.
- Veber, M. (2019). Why Not Persuade the Skeptic? A Critique of Unambitious Epistemology. *International Journal for the Study of Skepticism*: 1-25.
- Vogel, J. (1997). Skepticism and Foundationalism: a Reply to Williams, *Journal of Philosophical Research* 22: 11-28.
- White, R. (2006). Problems for Dogmatism, *Philosophical Studies* 131/3: 525-557.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Wittgenstein, L. (1953). *Philosophical Investigations*. Wiley-Blackwell.
- Wittgenstein, L. (1958). *The Blue and Brown Books*. Harper & Row.
- Wright, C. (1985). Facts and certainty. *Proceedings of the British Academy*, 71: 429–72.
- Wright, C. (2008). Internal—External: Doxastic Norms and the Defusing of Skeptical Paradox. *Journal of Philosophy* 105 (9): 501-517.
- Wright, C. (2010). The Illusion of Higher-Order Vagueness. In Richard Dietz & Sebastiano Moruzzi (eds.), *Cuts and Clouds. Vagueness, its Nature and its Logic*. Oxford University Press.
- Yablo. S. (1993). Paradox without Self-reference. *Analysis*. 53 (4): 251-252.

Zagzebski, L. (1994). The Inescapability of Gettier Problems. *Philosophical Quarterly* 44 (174): 65-73.

Zhao, H. (2018). Knowledge without Safety. *Synthese*: 1-18

Zimmerman, A. (2007). Against Relativism. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 133(3): 313-348.