



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.

Epistemic Fictionalism

Angela O'Sullivan

Doctorate in Philosophy (PhD)

The University of Edinburgh

May 2023

Declaration

I, Angela O'Sullivan, declare that (a) this thesis has been composed by me, and (b) that this work is my own, and (c) that the work has not been submitted for any other degree or professional qualification except as specified, and (d) this work has not published elsewhere, in whole or in part.

Signed:

Dated: 10th May 2023

Abstract

This thesis develops and defends epistemic fictionalism, according to which knowledge talk is metaphorical. One of the distinctive features of metaphor is that metaphorical sentences have multiple readings: a literal (or ‘face-value’) reading and at least one metaphorical (or ‘non-face-value’) reading. Typically, speakers who utter metaphorical sentences intend to communicate a content that corresponds to the metaphorical meaning. Epistemic fictionalism posits that, as is standard for metaphors, sentences of the form “S knows that P” admit of at least two different readings: a face-value (literal) reading, and a different non-face-value metaphorical reading. A face-value reading of sentences of the form “S knows that P” is that S is infallible with regards to P, and as such are always (or almost always) false at face-value. However, such sentences are not typically used to communicate their typically false face-value content, but some other, typically true metaphorical content.

The thesis applies the methods of function-first epistemology (Craig 1990, Hannon 2019a), taking as a starting point the question of what knowledge talk is used for. I argue that understanding knowledge talk as metaphorical coheres with the function that knowledge talk plays, in spite of all (or most) knowledge attributions being literally false: there are advantages to communicating the non-face-value, metaphorical content. Another key advantage of epistemic fictionalism is that it explains the intuitions surrounding sceptical arguments. On the one hand, sceptical infallibilism is compelling because it is the correct analysis of the face-value content of knowledge attributions. On the other hand, sceptical infallibilism is not threatening to everyday knowledge attributions because it does not threaten the truth of the metaphorical content that knowledge attributions are typically used to express. Overall, the main aim of the thesis is to motivate epistemic fictionalism as a serious contender worthy of further investigation and development.

Lay Summary

This thesis defends a novel and perhaps surprising proposal: that most talk about knowledge is metaphorical. On this view, sentences like, “I know that penguins eat fish”, “She knew the answers to all the questions on her maths test” and “I know that there is a castle in Edinburgh” are all metaphorical. These sentences can be interpreted either literally or metaphorically, and it is the non-literal metaphorical meaning that we typically try to get across when we claim to know something. I call this view ‘Epistemic Fictionalism’. *Epistemic*, which is philosophy’s term for “relating to knowledge”; and *Fictionalism*, to capture the sense in which these ‘related to knowledge’ sentences do not aim to communicate a literal truth, but rather, a metaphorical truth.

The thesis is thematically divided into three parts. The first part lays the groundwork by arguing that on a literal interpretation, sentences like “I know that there is a castle in Edinburgh” mean that I am infallible with regards to the question of whether there is a castle in Edinburgh. It also introduces the ‘function-first’ methodology used throughout this thesis, on which reflecting on the function or role of knowledge talk is central to investigating what the concept of knowledge is like. The second part develops my epistemic fictionalist view, specifying what the literal and non-literal readings of knowledge sentences are. In particular, I argue that knowledge talk functions in ways typical of metaphor, giving us good reason to think that it is in fact metaphorical. The final part of the thesis defends epistemic fictionalism from objections and compares it to other rival theories of knowledge. The main aim of the thesis is to motivate epistemic fictionalism as a serious contender worthy of further development.

In a nutshell, here are the two main arguments for thinking that knowledge talk is metaphorical. First, knowledge talk seems to fulfil the same function that metaphorical talk serves. Metaphors function to make otherwise complex information easier to process and think about. I argue that knowledge talk is like this, too: framing what it takes to be well informed in terms of being infallible confers several advantages, including allowing us to more easily reason about how well informed we are. Second, there is tension between what knowledge sentences mean, and what we use them to communicate. On the one hand, “I know that there is a castle in Edinburgh” intuitively means that there is no possibility of mistake. But we are fallible, and mistake is always possible: the castle might have been destroyed 5 minutes ago; my memory might be inaccurate, and I’m confusing the castle in Stirling with the Kirk in Edinburgh; or more radically, both the city of Edinburgh and its castle might be figments of a dream. On the other hand, surely when I say “I know that there is a castle in Edinburgh”, I do not intend to commit myself to the claim that it is impossible that the castle has just been destroyed. I just meant that I consider myself to be well informed on the question of whether Edinburgh has a castle – well enough informed to tell you about it before you visit, perhaps, but not so well informed that I think that I am infallible with regards to whether the castle has been destroyed 5 minutes ago. What explains this tension? My answer is that knowledge talk is metaphorical. Sentences like “I know that there is a castle in Edinburgh” have a literal meaning (i.e. I am infallible) which is typically false. But they are used to communicate a metaphorical meaning, which is often true (e.g. I am well informed). In this regard, knowledge talk is just like any other metaphorical talk.

Dedication

For my wonderful son Oscar, with all of my love.

Contents

Abstract.....	3
Lay Summary.....	4
Contents.....	6
Acknowledgements	9
Introduction.....	11
Chapter 1: Sceptical Infallibilism.....	15
1.1. The Threat of Scepticism	15
1.2. What Infallibilism Is	17
1.2.1. <i>Entailment Infallibilism</i>	18
1.2.2. <i>Ruling Out Infallibilism</i>	19
1.2.3. <i>Probability 1 Infallibilism</i>	20
1.2.4. <i>Certainty Infallibilism</i>	20
1.2.5. <i>We're not all fallibilists now</i>	22
1.3. Arguments for Infallibilism.....	22
1.3.1. <i>Epistemic Possibility and CKAs</i>	23
1.3.2. <i>Knowledge does not admit of degrees</i>	25
1.3.3. <i>Knowledge is Qualitatively Different to Non-knowledge</i>	29
1.3.4. <i>Knowledge is Closed under Known Entailment</i>	31
1.4. Conclusion.....	32
Chapter 2: Function-First Epistemology and the New Apraxia Objection to Sceptical Infallibilism	33
2.1. Function-First Epistemology	33
2.2. Anti-sceptical Function-First Epistemology	38
2.2.1. <i>The New Apraxia Objection to Scepticism</i>	38
2.2.2. <i>Double Bookkeeping</i>	40
2.3. Useful Error.....	42
2.3.1. <i>Strength of Epistemic Position</i>	43
2.3.2. <i>Communicated Content</i>	45
2.4. Conclusion.....	48
Chapter 3: Fiction and Function.....	49
3.1. Fictionalism in Philosophy.....	49
3.1.1. <i>Fictionalism: a first characterisation</i>	49

3.1.2. <i>Functional Fictional Talk</i>	53
3.1.3. <i>Putting the ‘fiction’ in ‘Fictionalism’</i>	55
3.1.4. <i>Five Criteria</i>	59
3.2. <i>Some Arguments for Fictionalism</i>	60
3.2.1. <i>Arguments from Ontology</i>	60
3.2.1. <i>Arguments from Function</i>	61
3.2.2. <i>Arguments from Linguistic Evidence</i>	61
3.4. <i>Chung’s Epistemic Fictionalism</i>	63
3.5. <i>Applying the 5 Desiderata</i>	65
3.5.1. <i>Representationally essential discourse</i>	66
3.5.2. <i>Procedurally essential discourse</i>	67
3.5.3. <i>Presentationally essential discourse</i>	68
3.6. <i>A Refinement: Chung’s Conceptual Metaphor Theory Epistemic Fictionalism</i>	70
3.6.1. <i>Objections</i>	71
3.7. <i>Conclusion</i>	73
Chapter 4: Epistemic Fictionalism, Function-First	74
4.1. <i>Function-First Epistemic Fictionalism: a proposal</i>	74
4.1.2. <i>Non-Face-Value and Face-Value Content</i>	75
4.1.3. <i>Metaphor Frames</i>	77
4.2. <i>Meeting the Fictionalist Criteria</i>	81
<i>i. A theory about sentences in a domain of discourse D</i>	81
<i>ii. Sentences in D would be infelicitous on a face-value reading</i>	81
<i>iii. Uttering sentences in D fulfils a useful function</i>	81
<i>iv. Uttering sentences in D fulfils a function in virtue of doing something other than communicating their face-value content</i>	81
<i>v. Applies the notion of fiction in the explanation of (iv)</i>	82
4.3. <i>The Problem of Scepticism</i>	86
4.4. <i>Advantages of Fictionalism</i>	94
4.4.1. <i>Epistemic Possibility and CKAs</i>	94
4.4.2. <i>Knowledge does not admit of degrees</i>	95
4.4.3. <i>Knowledge is Qualitatively Different to Non-knowledge</i>	95
4.4.4. <i>Knowledge is Closed under known entailment</i>	96
4.5. <i>Conclusion</i>	98
Chapter 5: Objections and Replies	99
5.1. <i>The Phenomenological objection: knowledge talk doesn’t feel/like fictional talk</i>	99

5.1.2. <i>Metaphor and Fictionalism</i>	103
5.1.3. <i>Pretence Force Metaphor Fictionalism</i>	105
5.1.4. <i>Non-Pretence Force Metaphor Fictionalism</i>	110
5.1.5. <i>Content Metaphor Fictionalism</i>	113
5.1.6. <i>Metaphor Fictionalism: knowledge talk shouldn't feel like fictional talk</i>	119
5.2. Dead Metaphors.....	120
5.3. What fixes the content of typically non-face-value talk?	123
5.4. Conclusion.....	126
Chapter 6: Epistemic Fictionalism and its Rivals	128
6.1 Hyperbole and Loose Talk.....	128
6.2 Knowledge as a Regulative Ideal.....	133
6.3 Epistemic Contextualism	138
6.3.1 <i>Linguistic Analogues</i>	139
6.3.2 <i>Everyday Context Shifts</i>	143
6.3.3 <i>A complication: is metaphor underpinned by contextualist mechanisms?</i>	147
6.4 Sceptical Pragmatic Invariantism	151
6.4.1 <i>Knowledge attributions in thought</i>	153
6.4.2 <i>The awareness objection</i>	156
6.5. Conclusion.....	163
Concluding Remarks	165
Bibliography	168

Acknowledgements

Many people have made this work possible. Above all, I am extremely grateful to Martin Smith for being an endlessly kind and supportive supervisor. I have learned so much from Martin's insightful and patient discussion of my work, and he has encouraged me at every turn to delve deeper into the questions that most excited me.

Thank you to my secondary supervisors, Matthew Chrisman and Aidan McGlynn, for their invaluable feedback and their guidance on navigating the academic landscape.

I am very grateful to my examiners Michael Hannon and Nick Treanor, whose careful and thoughtful engagement with my work has been a real privilege.

Thank you to the Scottish Graduate School of Arts and Humanities, for funding my research and for making me feel part of a Scotland-wide arts and humanities research scene.

I have gained immensely, both philosophically and personally, from all of the fellow graduate students and early career researchers that I have been fortunate enough to meet along the way. Particular thanks to Lilith Mace, whose friendship and boundless enthusiasm for talking epistemology have made these years, and this work, far better than they would otherwise have been. Thanks to Deryn Thomas, for being an unwaveringly supportive friend and for encouraging me to take seriously the practical import of philosophical work. For many fruitful and enjoyable discussions, thank you to Mark Ormeleas, Lilith Lee, Dario Mortini and Ross Patrizio. Thank you to all of the Scottish Epistemology Early Career Researchers (SEECRs), the most fantastic and friendly community from whom I have learned so much.

I am very grateful to everyone who has commented on earlier versions of the ideas in this thesis, especially members of the Varieties of Risk research group, Matthew Chrisman's Stammtisch group, the Edinburgh Graduate WiP and the University of Glasgow's COGITO epistemology research centre. I have learned from every event that I have presented at and attended.

Thank you to the City of Edinburgh, for being the most beautiful and special home over these last six years.

I would also like to thank the family and friends who have supported me throughout the past four years. Thanks to Heather Smith, who is always right. Thanks to Lizzie and Jason Rogers, who are fearless and who ask the best questions. Thanks to Mark Dimmock, who first introduced me to antirealism in an A level class 15 years ago, and who has been a source of guidance and spirited objections ever since. Thanks to Caro Part and Rana Fetit, whose friendship, encouragement and sense of humour pulled me through even the toughest days of balancing philosophy with new motherhood. Thanks to Pat and John Foley, for all of their help and support. Thanks to my parents, Wendy Yapp and Dermot O'Sullivan, who have always believed in me and encouraged me to pursue the things that I love.

Thank you especially to Luke Foley, without whom none of it would have been possible.

Finally, this thesis is dedicated to my wonderful son Oscar Foley, who was born at the end of the first year of this PhD and who has filled these last three years with so much love and joy.

Introduction

This thesis proposes and defends epistemic fictionalism, according to which knowledge talk is typically metaphorical.

One of the distinctive features of metaphor is that metaphorical sentences have multiple readings: a literal (or ‘face-value’) reading and at least one metaphorical (or ‘non-face-value’) reading. Typically, speakers who utter metaphorical sentences intend to communicate a content that corresponds to the metaphorical meaning. For example, there are two readings of Romeo’s metaphorical sentence “Juliet is the sun”: a literal reading, on which the sentence communicates that Juliet is a flaming ball of plasma around which planet Earth rotates; and a metaphorical reading, on which the sentence communicates something along the lines of Juliet is unique, radiant, beautiful, the object around which Romeo’s thoughts and dreams centre. It is this second, metaphorical content that Romeo intends to communicate, and which we as co-operative hearers understand him to mean. Metaphorical talk is not confined to literature. Consider this everyday metaphor: “we’ve come so far”. This metaphor is so ubiquitous that it might not initially register as metaphorical at all. But it too admits of two readings: a face-value reading, on which it communicates that we have physically travelled across a long distance, perhaps several miles; and a non-face-value metaphorical meaning, on which it communicates that we have made significant progress towards a shared aim, perhaps in our personal or professional lives. When this sentence is used metaphorically in context, co-operative hearers readily understand that the speaker intends to communicate something about their progress towards a shared aim, rather than something about physical miles travelled.

To say that knowledge talk is metaphorical is to say that, just like “Juliet is the sun” and “we’ve come so far”, sentences of the form “S knows that P” admit of at least two different readings: a face-value literal reading and a non-face-value metaphorical reading. Just as with typical metaphors, it is the face-value reading that speakers typically intend to communicate, and which co-operative interlocutors understand them to be communicating. The general claim that knowledge is metaphorical is compatible with a range of different theories on what the face-value and non-face-value readings of sentences of the form “S knows that P”. In this thesis, I argue that the face-value reading is that posited by the sceptical infallibilist: that S has eliminated every possibility in which $\sim P$, that S’s belief that P is based on evidence that entails that P, that the probability of P on S’s evidence is 1, and so on. However, sentences of the form “S knows that P” are not typically used to communicate this face-value content. I argue that the non-face-value metaphorical content that we do use them to communicate is tied to the function that knowledge plays in everyday life: perhaps, that S is a reliable informant with regards to P, that S may stop inquiring into whether P, that S can provide assurance as to P, and so on. Crucially, it is this content that sentences of the form “S knows that P” typically communicate.

I argue that understanding knowledge talk as metaphorical makes sense of otherwise puzzling features of knowledge talk. One of these puzzles is the problem posed by sceptical arguments, and in particular by sceptical infallibilist arguments. Sceptical infallibilism is the view that the standards for knowledge are maximally and unattainably demanding, and therefore most sentences of the form “S knows that P” are false. Sceptical infallibilism seems untenable because it commits everyday speakers to widespread and systematic error. It is difficult to understand how we could be so confused regarding large swathes of our talk, and difficult to account for why such a demanding concept would ever have emerged and remained in circulation. Yet, sceptical

infallibilism is explanatorily powerful on a range of issues, and its arguments elude easy refutation. One of the central contributions of this thesis is a solution to this problem. On epistemic fictionalism, the sceptical infallibilist is correct about the face-value reading of sentences of the form “S knows that P”. The reason that sceptical infallibilism seems compelling and eludes refutation is that it is, literally, true. However, on my picture the literal truth of sceptical infallibilism is not disturbing or threatening, and does not commit us to widespread error. This is because sentences of the form “S knows that P” are not used to express their (false) face-value content – but some other, non-face-value content, one that is typically true. Epistemic fictionalism, then, diagnoses both the appeal of and resistance to sceptical arguments.

Throughout this thesis, I apply the methods of function-first epistemology (Craig 1990, Hannon 2019a), according to which reflecting on the function of a concept or domain of discourse illuminates what that concept is like, for example, what intension or extension it has. The arguments in this thesis are guided by reflection on the function of knowledge in multiple, but complementary, ways. At one level, hypotheses about the function of the concept of knowledge directly inform what metaphorical content I take sentences of the form “S knows that P” to communicate. But the role of function in my thesis goes beyond this. I argue that metaphorical talk in general serves a useful function: we speak metaphorically when there is some advantage to communicating a particular content via metaphor, as opposed to communicating it at face-value. I then argue that knowledge talk also serves this function, giving us positive reason to think that knowledge talk is metaphorical. That is to say, I argue that there are key advantages to metaphorically framing what it takes to be a reliable informant in terms of being infallible. One of the main aims of this thesis is to showcase how appealing to function can inform an anti-realist picture of knowledge on which knowledge talk is both useful, and typically false on a face-value reading.

Along the way, I also contribute towards live debates surrounding fictionalism in general. For example, I argue that fictionalism is better motivated when it includes an account of how the notion of fiction enables the target discourse to fulfil its function. I also argue that metaphor fictionalism is well placed to address the phenomenological objection to fictionalism. Finally, I use reflections on the function of a discourse domain to answer the question of how a typically non-face-value discourse domain derives its face-value content. The hope is that these points not only serve to make plausible epistemic fictionalism, but fictionalism in other areas also.

The thesis is comprised of six chapters which (with apologies to one of my examiners, who does not enjoy reading chapter summaries) I now briefly outline.

Chapter 1 is a defence of sceptical infallibilism. My epistemic fictionalist view emerges as a way of handling the (face-value) truth of sceptical infallibilism, and so this chapter lays out the motivation for epistemic fictionalism. First it outlines four popular ways of spelling out infallibilism about knowledge and argues that all of them generate scepticism. Then, it discusses some key arguments in favour of sceptical infallibilism: explaining why concessive knowledge attributions are infelicitous; explaining the qualitative difference between knowledge and non-knowledge; and explaining the intuition that knowledge is closed under known entailment. It also spells out the sense in which scepticism is threatening, and why epistemologists have traditionally sought to avoid sceptical infallibilism.

Chapter 2 introduces the function-first methodology which guides this thesis, and responds to a function-first objection to sceptical infallibilism. Applying function-first epistemology involves

investigating a concept by reflecting on what function the concept serves – on what we could not do, or could not do so easily, without this concept. Several function-first epistemologists have argued that sceptical infallibilism is incompatible with the point of knowledge, since no function is served by a concept that systematically fails to apply. I defend sceptical infallibilism from this objection, arguing that a concept or discourse domain may be functional in spite of typically being used in false sentences. In doing so, I motivate the broader class of anti-realist views to which fictionalism belongs.

Chapter 3 introduces fictionalism as a general theory and defends five criteria that constrain a well-motivated fictionalist view, as well as highlighting some of the arguments generally employed to motivate fictionalism about a discourse domain. It then examines in detail the only other iteration of epistemic fictionalism currently on the market: Julianne Chung's (2018, 2020) epistemic fictionalist theory. I argue that despite its many virtues, Chung's fictionalism does not meet one of the five criteria for a well-motivated fictionalist theory: it does not account for the function of non-face-value knowledge talk. Examining Chung's view will serve as a springboard from which to develop the version of fictionalism that I prefer.

Chapter 4 develops my positive proposal, a 'function-first' epistemic fictionalism underpinned by the mechanism of metaphor frames. On this picture of epistemic fictionalism, knowledge talk is metaphorical talk, and is typically used to express a metaphorical content. This metaphorical content relates to the roles that the concept of knowledge serves. Having outlined my positive proposal, I highlight how it meets the five criteria for a well-motivated fictionalist theory. Then, I put it to work in resolving the problem presented by sceptical infallibilism. I argue that the face-value reading of sentences of the form "S knows that P" is that posited by the sceptical infallibilist, which explains why sceptical infallibilism is compelling: it is, literally, true. Yet, epistemic fictionalism avoids the problems commonly associated with sceptical infallibilism: sentences of the form "S knows that P" typically communicate a true content, and so everyday speakers are not committed to widespread error. Finally, I explain how epistemic fictionalism retains the advantages of sceptical infallibilism outlined in Chapter 1.

Chapter 5 moves on to defending epistemic fictionalism by addressing three pressing objections: that knowledge talk does not feel anything like a fiction; that the knowledge metaphor is so common that its metaphorical meaning has 'died' and become part of its face-value meaning so that knowledge talk is actually better thought of as face-value talk; and that epistemic fictionalism struggles to explain where knowledge talk derives its face-value content. In the course of answering these objections, I demonstrate two distinct ways of fleshing out epistemic fictionalism as either assertoric (content) fictionalism or non-assertoric (force) fictionalism.

Finally, Chapter 6 continues the defence of epistemic fictionalism by comparing it to rival views. First, I compare epistemic fictionalism to three other anti-realist views: hyperbole, loose talk and regulative ideals. I argue that among other advantages, epistemic fictionalism better captures the function of knowledge talk. Then, I compare epistemic fictionalism to what I consider its main rivals: epistemic contextualism and sceptical pragmatic invariantism, both of which also posit that the sceptical infallibilist is in an important sense right about (at least one of) the meaning(s) of sentences of the form "S knows that P". These views appear at first pass to generate the same results as epistemic fictionalism, without the burden of positing non-face-value mechanisms. In response, I argue that there are 'counterpart' ways of specifying epistemic fictionalism for both epistemic contextualism and sceptical pragmatic invariantism, and that these counterparts retain

the advantages of epistemic contextualism and sceptical pragmatic invariantism respectively, while additionally resolving objections to the face-value version of each view.

The main aim of this thesis is to motivate epistemic fictionalism as a serious contender worthy of further investigation and development. Above all, I want to show that our everyday concept of knowledge could be functional, even if fictional – indeed, that it might be functional *in virtue of* being fictional. A metaphorical concept of knowledge would be no cause for despair. Quite the opposite, if knowledge talk turns out to be metaphorical then we have every reason to carry on using our functional, fictional, infallibilist concept of knowledge.

Chapter 1: Sceptical Infallibilism

This thesis defends a view on which sceptical infallibilism is, in one sense, correct: sceptical infallibilism captures the content that sentences of the form “S knows that P” communicate if taken at face-value. Infallibilism about knowledge is an unpopular view among contemporary epistemologists. The main reason for this is that infallibilism is strongly associated with scepticism about knowledge, understood as the claim that all (or most) of our uttered sentences of the form “S knows that P” are false. Spelling out precisely how to understand infallibilism is not straightforward. But an initial, rough, characterisation is this: infallibilism holds that the standards for knowledge are maximally demanding. Because accepting infallibilism means accepting that the standards for knowledge are maximally demanding, epistemologists worry that accepting infallibilism about knowledge means accepting that scepticism is true.

Yet, despite this, infallibilism is explanatorily powerful. In addition to capturing the intuitive sense in which knowing that P means that one could not be wrong about P, infallibilism explains otherwise puzzling data: why concessive knowledge attributions are infelicitous; why knowledge seems to be qualitatively different to non-knowledge; and why knowledge can be extended by competent deduction. One of the central themes of this thesis is that (on one level) sceptical infallibilism about knowledge is correct, but yet sceptical infallibilism is not threatening. The chapter proceeds as follows. First, I outline the type of scepticism at issue in the debate around infallibilism (§1.1.). Then, I canvass various formulations of infallibilism and argue that each one generates scepticism (§1.2.). Finally, I lay out the arguments in favour of infallibilism, with a particular focus on the infelicity of concessive knowledge attributions (§1.3.).

1.1. The Threat of Scepticism

Scepticism is the view that most sentences of the form “S knows that P” are false, and perhaps, that most sentences of the form “S is justified in believing that P” are also false. That a theory of knowledge entails scepticism is generally thought to count against it. Jim Stone writes that “a subtext of virtually every theory of knowledge has been to show how knowledge is possible or, at the least, to avoid an account that delivers us unto the skeptic” (2000:527). Indeed, Scepticism finds few adherents among contemporary epistemologists¹, who tend to regard the truth of the sceptic’s conclusion as “intellectually disastrous” (Pritchard 2008:39), “depressing and impractical” (Hannon 2019b:143) and “wholly unacceptable” (Williams, 2004:457). If infallibilism generates the conclusion that we know nothing or very little, then this is generally considered a good reason to reject it.

Following Hilary Kornblith (2000:25-6), we can distinguish between two broad types of scepticism. ‘High-standards’ sceptics hold that the standards for knowledge are so demanding as to go unmet in almost every context in which knowledge attributions are typically made. One way of being a high-standards sceptic is to hold that knowing a proposition P (e.g. ‘I have hands’) requires ruling out every possibility in which $\sim P$, including remote possibilities and possibilities

¹ See Stroud (1984), Stone (2000), Schaffer (2004), Cappelen (2005), Davis (2007) Dodd (2011), Bonjour (2010) Stoutenburg (2021) and Rinard (2021) for scepticism-friendly exceptions.

that no obtainable piece of evidence would suffice to rule out (e.g. ‘I am a handless brain in a vat being systematically deceived’). Another way of being a high-standards sceptic is to think that knowing a proposition *P* requires that one’s justification for (or evidence for, or reason to believe) *P* entails the truth of *P*. However the details are fleshed out, on high-standards scepticism most of our knowledge attributions turn out false. However, the high-standards sceptic allows that although no belief amounts to knowledge, some beliefs nevertheless score more highly than others on other epistemic measures. For example, the high-standards sceptic can allow that some beliefs are more justified, more warranted, or better supported by evidence when compared to other beliefs. In contrast to high-standards scepticism, ‘full-blooded’ (Kornblith 2000:27) or ‘evidential’ (Rinard 2021:2) scepticism holds that in addition to lacking knowledge, we also “have no degree of justification whatever for [our] claims about the external world. None.” (Kornblith 2000: 26). That is to say, for the full-blooded sceptic, no belief is epistemically better off than any other – no belief is more justified, enjoys more warrant, or is better evidenced than any other.

This thesis is concerned with high-standards scepticism. Some philosophers find high-standards scepticism to be uninteresting and unthreatening. Hannon, for example, writes that “as soon as we admit that some beliefs are epistemically better than others, scepticism loses its teeth” (2019b:154); Butchvarov writes that “the interesting sort of skeptic is not the one who denies that we know what we think we know but the one who denies that we have good (or even any) reasons for believing what we think we are justified in believing” (1970:59); and Kornblith writes that high-standards scepticism is a “wholly trivial and uninteresting position” (2000:27).

I agree that full-blooded scepticism is the more threatening of the two positions, since it threatens not only out intuitive knowledge attributions, but epistemic evaluation as a whole: our intuitive judgements that a given belief is supported by evidence, or likely to be true, or sufficiently motivated given the context, and so on. I do however disagree that high-standards scepticism is harmless and unthreatening. For high-standards scepticism embroils everyday users of knowledge discourse in widespread and systematic error. If most of our knowledge attributions go unmet, then large swathes of the sentences that we take to be true are actually false. Either we are mistaken about the meaning of the word ‘knows’, and thus systematically deceived about the meaning of perfectly ordinary sentences; or we are mistaken about our epistemic position, mistakenly believing ourselves to be better epistemically positioned than we in fact are. Neither consequence is desirable. Furthermore, as Hannon (2015) points out, the problem isn’t limited to the English language term ‘knows’: equivalent concepts feature in every language, and are used in comparable ways. The truth of high-standards scepticism, then, would not only have as a consequence that users of the English word ‘knows’ and its cognates are mistaken about meaning or strength of epistemic position, but more generally that almost all language users are so mistaken². Finally, many epistemologists have argued that knowledge is the norm of assertion³ and practical reasoning⁴. On these views, one may assert that *P* and act on *P* only one knows that *P*. If this is right, then the truth of scepticism would mean that the norms of assertion and practical

² In supporting this point, Hannon cites Goddard (2010) and Haspelmath and Tadmor (2009)

³ See for example Unger (1975), Williamson (2000), DeRose (2002), Reynolds (2002), Adler (2002), Schaffer (2008), Jackson (2012), Turri (2016) and Mentha (2016).

⁴ Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), and Hawthorne and Stanley (2008)

reasoning are systematically violated. Scepticism would threaten not only our knowledge attributions, but assertions in general, and practical reasoning also. Thus high-standards scepticism is threatening in the sense that it threatens to commit everyday language users to widespread error and norm violation.

Finally, both friends and foes of infallibilism tend to agree that “fallibilism appears to have carried the day in epistemology” (Reed 2013:410) due to its strong association with scepticism. Critics object that infallibilism entails scepticism; scepticism is to be rejected; and therefore, infallibilism is also to be rejected. Famously, Cohen (1988:91) writes that “the acceptance of fallibilism in epistemology is virtually universal [and] derives from the widely held view that what we seek in constructing a theory of knowledge is an account that squares with our strong intuition that we know many things”. This diagnosis is close to consensual. More recently, it is explicitly echoed by Dodd (2011:669), who writes that “most contemporary epistemologists will want to reject infallibilism, because they think infallibilism leads to scepticism, a fact which they think provides a *reductio* of infallibilism.”; by Brown (2018:2) who says that “Fallibilism was the default view in epistemology for much of the twentieth century. For, many thought that endorsing fallibilism was the only way to avoid scepticism about knowledge”; and by Fratantonio, who writes that “infallibilism about knowledge has a bad reputation for it is often thought to easily lead to scepticism” (2021:1307). That scepticism is widely regarded as the cause of the widespread adoption of fallibilism is testament to how threatening scepticism is: scepticism is so threatening that any theory of knowledge that appears to entail it is to be rejected. But if infallibilism about knowledge does indeed entail scepticism, it entails *high-standards scepticism*, not full-blooded scepticism. That the standards for knowledge are maximally demanding carries no *prima facie* entailment as to whether any of our beliefs are epistemically better off than any others; for example, on whether any of our beliefs are better supported by evidence, or more justified. Say that infallibilism about knowledge is correct, and the standards for knowledge are maximally demanding: it may well still be the case that beliefs can be more or less justified, or can approximate knowledge to a closer degree than others, and so on. High-standards scepticism, then, is a threatening view: so threatening that many epistemologists have taken it as sufficient cause to wholesale reject sceptical infallibilism.

1.2. What Infallibilism Is

Despite the almost universal endorsement of fallibilism (or, as Reed 2002:143 suggests, “perhaps because of it”), there is no consensus on how best to characterise infallibilism and fallibilism. Most epistemologists seek to preserve the intuition that the infallibilist’s concept of knowledge is based on evidence that entails the truth of what is known, or is incompatible with the possibility of error. In contrast, fallibilism is typically described as the denial of this thesis. Fallible knowledge can be acquired even in absence of evidence that guarantees the truth of what is known, or even if there are un-eliminated possibilities of error.

Here are four mainstream formulations of infallibilism; fallibilism is the denial of the infallibilist claim:

- ENTAILMENT:** If S knows that P then S's knowledge that P is based on evidence that entails P.
- RULING OUT:** If S knows that P then S has ruled out every possibility in which $\sim P$
- PROBABILITY 1:** If S knows that P, then the probability that P on S's evidence is 1
- CERTAINTY:** If S knows that P, then S is certain that P⁵

I now briefly discuss each type of infallibilism, arguing that each one, in at least one natural iteration, generates scepticism.

1.2.1. Entailment Infallibilism

Many epistemologists formulate the distinction between fallibilism and infallibilism as the distinction between whether knowledge requires truth-entailing evidence. In a nutshell, the idea is that for the infallibilist, knowledge that P requires that one's evidence for P entail the truth of P. For example, Feldman characterises fallibilism as the view that "it is possible for S to know that P even if S does not have logically conclusive evidence to justify believing that P"; Stanley (2005:127) defines fallibilism as "the doctrine that someone can know that p, even though their evidence for p is logically consistent with the truth of not-P"; Littlejohn (2011:603) likewise writes that "[f]allibilists think it's possible for you to know that p is true even if your evidence for believing p does not entail that p is true"; and Brown (2018:41) characterises infallibilism as the view that "one can know that p only if one's evidence for p entails that P".⁶

Key to entailment infallibilism, then, is the notion of evidence. Depending on how we understand the requirement that S's evidence entails P, we'll end up with radically different conceptions of infallibilism. If we understand evidence along internalist lines to capture S's subjective experiences, then the entailment infallibilist requirement is that S be able to distinguish the experience of P from the experience of $\sim P$, for any $\sim P$ possibility. If however we understand evidence along the lines suggested by Williamson (2000), then the entailment infallibilist requirement is the requirement that knowing that P demands that P must be entailed by S's evidence.

At least in its internalist guise, entailment infallibilism straightforwardly generates scepticism. To know that P, S's evidence must entail that P: but for any proposition P, it is possible to generate a proposition that is compatible with S's evidence and compatible with some $\sim P$ possibility. Sceptical hypotheses are the paradigmatic example of this phenomenon. On one way of

⁵ 1 is the formulation used by Lewis (1996) and Pritchard (2005). For 2, see Stanley (2005), Cohen (1988), Dougherty (2011), Neta (2011) and Brown (2018) among many others. For 3, see Williamson (2000), Littlejohn (2008), Dodd (2011), and Fratantonio (2021). For 4, see Descartes (1641) and Unger (1971, 1975).

⁶ Other formulations of fallibilism along these lines include Cohen (1988), Brueckner (2005), Dutant (2007), Dougherty (2011), Neta (2011), Reed (2012) and Brown (2018), among many others.

constructing sceptical arguments, the sceptic raises a possibility which is compatible with S's evidence (for example, that she is a handless brain in a vat being stimulated into having experiences as if of the external world), but incompatible with the proposition that she claims to know (for example, that S has hands). If our evidence for P does not entail the truth that P then, according to entailment infallibilism, we fail to know that P – and scepticism ensues⁷.

As discussed, this way of understanding entailment infallibilism hinges on its internalist reading. If we understand evidence as being comprised of the propositions that S knows, then entailment infallibilism need not yield scepticism. For on these 'knowledge views' of evidence, what counts as one's evidence is not one's subjective experience as if of sitting at a desk writing, but *that* one is sitting at the desk writing. It is thus not challenging to acquire evidence that entails the proposition that one knows. If one can possess lots of infallibilist knowledge, then scepticism is avoided. Indeed, adopting the knowledge view of evidence on all four of these versions of infallibilism avoids scepticism. I return to the question of non-sceptical externalist infallibilism at the end of this chapter.

1.2.2. Ruling Out Infallibilism

Epistemologists who draw the divide in the manner formulated by ruling out infallibilism include Lewis (1996) and Pritchard (2005). Pritchard defines infallibilism as “the thesis that, roughly, for an agent to know a proposition that agent must be able to eliminate *all* error-possibilities associated with that proposition” (2005:218). And Lewis defined his infallibilist starting point as the thesis that “S knows proposition P iff P holds in every possibility left uneliminated by S's evidence” (1996:551) to contrast with fallibilism, which legitimises statements such as “[h]e knows, yet he has not eliminated all possibilities of error” (1996:550).

What does it mean to say that S has ruled out every possibility in which $\sim P$? First, there are two ways of reading the requirement that a possibility be 'ruled out'. One is the demand for reflective ruling out, so that S is reflectively aware of which propositions she has eliminated. This principle is a version of the now unpopular internalist demand that knowledge be 'luminous' (Williamson, 2000), 'transparent' (Pritchard, 2001) or 'directly before consciousness' (Fumerton, 1990:64). To say that knowledge is luminous or transparent is to say that an agent is always able to discern, by reflection alone, that she knows that P. This version of ruling out infallibilism straightforwardly generates scepticism. On the plausible assumption that it is always possible to construct a possibility which S cannot eliminate but which is incompatible with the would-be-known proposition, then one is never in a position to eliminate every possibility. Returning to sceptical scenarios, according to the ruling out infallibilist, that S cannot rule out the possibility that she is a brain in a vat means that S cannot acquire any knowledge about the external world⁸.

In contrast, the non-reflective version holds that ruling out is a function of what S knows, irrespective of whether this knowledge is accessible to S. On this way of spelling out the view, that S must rule out all possibilities that $\sim P$ in order to know that P is trivial: to say that there are

⁷ This iteration of infallibilism is what underpins sceptical arguments from underdetermination (Smith 2022, Brueckner 1994).

⁸ This way of thinking about infallibilism underpins closure-based sceptical arguments (see §1.3.4.).

possibilities left uneliminated by S's evidence is just to say that these possibilities are compatible with S's knowledge. Put differently, if S doesn't know that P, then $\sim P$ is not ruled out for S.

1.2.3. Probability 1 Infallibilism

Probability 1 infallibilism has recently received revived attention. Probability 1 infallibilism characterises infallibilism as the thesis that “[i]f a subject S knows P, then the epistemic probability of P for S is 1 (PSP = 1)” (Dodd 2011:665). Probability 1 infallibilism has been defended by Williamson (2000), Littlejohn (2008), Dodd (2011), and Fratantonio (2021), of which only Dodd (2011) takes probability 1 infallibilism to generate scepticism.

Dodd understands the relevant notion of epistemic probability as being relative to S's epistemic situation, which he cashes out in terms of S's evidence⁹. If, relative to S's evidence, there is an epistemic possibility that $\sim P$, then the probability of P is less than 1. Therefore for Probability 1 infallibilism if, relative to S's evidence, there is an epistemic possibility that $\sim P$, then S does not know that P. The “crucial assumption” (Dodd 2011:667) in the argument is that epistemic possibility is a function of the agent's knowledge: “P is epistemically possible for S only if S doesn't know the contradictory of P to be true”. So: according to the Probability 1 infallibilist, for S to know that P requires that P not be epistemically possible for S; which in turn requires that S know $\sim P$.

Here is how Probability 1 infallibilism generates scepticism: for almost all of the propositions that we take ourselves to know, unless P is entailed by S's evidence, then there will be some non-0 probability that $\sim P$ given S's evidence. For example, unless S's evidence that she has hands entails that she has hands, then there will be some probability, given S's evidence, that S does not have hands. The probability for S that she has hands will be, given her evidence, high (99.99%, say) but nevertheless it will be less than 1. Given this, on Probability 1 infallibilism, we don't know many of the propositions that we take ourselves to know.

1.2.4. Certainty Infallibilism

The certainty requirement attributed to Descartes (1641), Unger (1971, 1975) and, to an extent, Moore (1959) is standardly understood as an infallibility requirement. There are two ways of understanding ‘certainty’: *psychological* certainty and *epistemic* certainty. Psychological certainty is supposed to capture one's internal subjective feeling of maximal confidence in the truth of their belief. In contrast, epistemic certainty is an epistemic property of beliefs. S is epistemically certain that P if S knows (or is in a position to know) that P, and if S's belief that P is formed “on the basis of evidence that gives one the highest degree of justification for one's belief” (Stanley 2008:35). Epistemic certainty is typically accompanied by psychological certainty, but need not be.

On Psychological Certainty Infallibilism, if S knows that P then S feels maximally confident that her belief that P is true. Psychological Certainty Infallibilism does not entail scepticism, or at a

⁹ More carefully, Dodd is open to understanding ‘epistemic situations’ in terms of evidence.

minimum is does not logically entail scepticism. For whether it is possible to feel maximally confident in one's beliefs is an empirical question to be answered through investigations into our attitudes and dispositions, and it may well turn out that such an attitude is possible. Unsurprisingly, no contemporary epistemologist defends Psychological Certainty Infallibilism (as far as I am aware). After all, psychological certainty just does not seem to reliably track strength of epistemic position, undermining the rationale for such a requirement on knowledge. For example, many think that the student in Radford's (1966) case does not fail to know the date of the battle of Hastings, even if during the test he doubts his own memory. And although these are not counter-examples to the psychological certainty requirement as I have formulated it, a tea-leaf's reader or blackjack player's maximal psychological confidence is clearly no gauge at all as to whether their belief amounts to knowledge, putting further pressure on the rationale for such a requirement¹⁰.

Unger (1971, 1975) uses psychological certainty as the basis of an epistemic certainty requirement on knowledge. Unger argues that knowledge licences (psychological) certainty, so that if S knows that P then A is entitled to be (psychologically) certain regarding P. This is because for Unger 'certainty' is an absolute term, like 'empty' or 'flat'. If a room is empty then it is not at all populated, if a surface is flat then it is not at all flat, and if a proposition is certain then it is not at all doubtful. Further, certainty is the attitude in which one is prepared to gamble on P at any stakes and, more importantly, to dismiss any future possibility in which the known proposition is false. However, Unger goes on to argue that certainty is never warranted: so long as there is a live possibility of error, then it is possible for a proposition to be less doubtful.

Unger's argument proceeds as follows:

U1. If I know that P, I am entitled to be certain that P

U2. I am not entitled to be certain that P

UC. Therefore, I don't know that P.

Unger-style infallibilism, then, is ultimately a version of Ruling Out Infallibilism, according to which if there are uneliminated error possibilities, then one is not warranted in feeling certain that P. Since if error possibilities are never completely eliminated, certainty is never warranted, and scepticism follows.

¹⁰ Nevertheless, I think that psychological certainty infallibilism is the right reading of Descartes' requirement on knowledge, combined with Descartes' view that (a) rational beings can reliably detect error possibilities and (b) contemplating error possibilities destroys psychological certainty. (a) captures the rationalist ideal that if we are sufficiently rational, we can clearly and distinctly see truth and avoid error. For Descartes, psychological doubt acts as a gauge in rational agents, detecting the possibility of error and alerting us to the absence of knowledge. See Cuning (2009) for a defence of the claim that Descartes understood the sufficiently rational agent as being able to identify, and so avoid, error.

1.2.5. *We're not all fallibilists now*

I have tried to sketch out how various formulations of infallibilism generate scepticism. All of these views try to capture the intuitive everyday sense in which knowing that *P* means that *S could not be wrong about P*. As discussed, many epistemologists take infallibilism's sceptical implications to speak decisively against it. Nevertheless, it is important not to overstate the extent to which "we are all fallibilists now" (Siegel 1997:164). As McGlynn rightly points out, although in 1988 Cohen could confidently write about the universal acceptance of fallibilism, nowadays "things look decidedly different" (2019). Infallibilism has indeed enjoyed somewhat of a resurgence since 1988, in large part owed to the introduction of approaches which seek to divorce infallibilism from scepticism. The two most popular examples include epistemic contextualism, with its promise of avoiding both "the rock of fallibilism and the whirlpool of scepticism" (Lewis 1996:550); and more recently, evidential externalist views which employ what Jessica Brown refers to as a "generous conception of evidence" (2018:1, 3-4) on which knowledge that *P* is based on factive evidence that *P*. The question of whether these new theories entail scepticism takes centre stage in the revived fallibilism-infallibilism debate, with fallibilists insisting that the new infallibilist theories do not avoid scepticism – or at least, that they avoid it only at unbearable cost (Brown 2018), and defenders of new versions of infallibilism arguing that their infallibilism does not also commit them to scepticism¹¹.

A discussion of externalism and internalism goes beyond the scope of this thesis, which focuses on non-externalist versions of infallibilism. I welcome the externalist-driven revival of infallibilism. It has breathed new life into what had been previously regarded as a closed debate, and has succeeded in making externalist anti-sceptical infallibilism a reasonably uncontroversial position. But sceptical infallibilism remains as unpopular as it ever was, and sceptical consequences are still viewed by anti-sceptical (externalist) infallibilists and fallibilists alike as a serious reason to reject sceptical infallibilism. It therefore is surely of interest, for both anti-sceptical externalist infallibilists and fallibilists, if it can be shown that sceptical infallibilism need not be threatening. That is the aim of this thesis. I will not, therefore, have much to say about anti-sceptical externalist infallibilism; my main interlocutors are anti-sceptical fallibilists (for example, chapter 5 addresses loose talk and exaggeration views; chapter 6 addresses contextualism and non-sceptical invariantism). Nevertheless, I hope that the arguments raised in this thesis might be of interest to infallibilists of all stripes, anti-sceptical ones included.

1.3. Arguments for Infallibilism

I now turn to a defence of sceptical infallibilism (henceforth, just 'infallibilism'), by arguing that infallibilism better explains three otherwise puzzling features of knowledge: why concessive knowledge attributions of the form "I know that it's raining, but it might not be raining" sound infelicitous; why knowledge is qualitatively different from nearly-knowledge; and why knowledge can be extended via the principle of closure. More attention is given to what I consider to be the most persuasive argument for infallibilism: that fallibilism is incompatible with standard accounts

¹¹ See for example Williamson (2000), Dutant (2007), Booth (2017), Venturinha (2020), Fratantonio (2021).

of epistemic modality, and hence the concessive knowledge attributions licenced by fallibilism sound infelicitous. I consider the response that knowledge admits of degrees, before more briefly discussing the two other features of knowledge that infallibilism is well placed to explain.

1.3.1. Epistemic Possibility and CKAs

To my mind, the most persuasive argument for infallibilism is that fallibilism is incompatible with standard accounts of epistemic modality (Dodd 2011, Reed 2013, Anderson 2014, Climenhaga 2021). On standard accounts of epistemic modality, operators ‘might’ and ‘must’ quantify over epistemic possibilities, where epistemic possibilities are possibilities compatible with an agent’s knowledge. These operators mark the necessity (‘must’) or possibility (‘might’) of a proposition, relative to some body of knowledge¹². The standard account of epistemic modality holds that if “P must be the case” for S, then P is entailed by S’s evidence; and if “P might be the case” for S then P is not entailed by S’s evidence. If P might be the case for S, then P is compatible with S’s knowledge. Operators quantify over epistemic possibilities, which are the possibilities compatible with an agent’s knowledge. If ‘it might be that P’ is true for S then it is epistemically possible for S that P. If ‘it might be that P’ is false for S, then it is epistemically impossible for S that P. This can be put more formally in the following principles:

Epistemic Possibility (EP): $\sim P$ is epistemically possible for S only if S does not know that P.

And by contraposition:

Epistemic Impossibility (EI): If S knows that P, then $\sim P$ is not epistemically possible for S.

These two logically entailed principles are widely accepted. Explicit formulations include:

Whenever a speaker S does or can truly assert that ‘It’s possible that P is false’, S does not know that P. (DeRose, 1991:596)¹³

It is possible that P for S at t (There is a chance that P for S at t) iff P is consistent with what S knows at t. (Hawthorne, 2004:24)

It is possible for A that P is true if and only if what A knows does not, in a manner that is obvious to A, entail not-P. (Stanley, 2005:128)

Where it is not explicitly stated, this reading is often implicit. It is also highly intuitive, an intuition supported by the infelicity of so-called concessive knowledge attributions, or CKAs (Rysiew, 2001). These are expressions of the form ‘I know P, but it might be that $\sim P$ ’ or ‘I know P, but it

¹² Kratzer (1977, 1981, 1991), based on Kripke (1963) and Hintikka (1962). See also Copeland (2002).

¹³ DeRose calls this ‘Moore’s Principle’ and takes it to be a constraint on EP. The full definition is: It is possible that P is true if and only if (1) no member of the relevant community knows that P is false, and (2) there is no relevant way by which members of the relevant community can come to know that P is false.

is possible that Q' where Q entails $\sim P$ in a manner obvious to S. The following sentences, for example, are infelicitous and even a charitable interlocutor would struggle to interpret the speaker's intended meaning:

“I know that Alana committed the murder, but it's possible that she didn't”

“It's possible that Belle is in her office, but I know that Belle's in the garden”

The standard account of epistemic modality given above easily explains why such CKAs are infelicitous: they contradict EP¹⁴. The 'possible' here is naturally understood as expressing epistemic, rather than logical or metaphysical possibility. If it's possible for me that Belle is in her office, then by EP I cannot know a proposition incompatible with the proposition that Belle is in her office, including the proposition that Belle is in the garden. Furthermore, this way of understanding epistemic modals coheres with the function that employing such modals serves in everyday life. As Dodd (2011:668) puts it, “though the expression ‘epistemic possibility’ is a philosopher's invention, the concept of epistemic possibility plays an important role in everyday thinking”, namely, to describe one's epistemic situation. Sentences such as “Belle might be in her office” describe something about the speaker and her epistemic situation, rather than something about Belle. They serve the function of explicating what possibilities, relative to the agent's knowledge, are left open.

Fallibilism is incompatible with the two epistemic possibility principles on all four of the characterisations outlined above. Ruling Out fallibilism most obviously conflicts with EI, for it says that S can know that P, even though S has not ruled out some possibilities in which $\sim P$. But to say that S has not ruled out some possibilities in which $\sim P$ is just to say that $\sim P$ is epistemically possible for S. That S can know that P despite some $\sim P$ possibilities being epistemically possible for S is in contradiction with EI. To endorse Ruling Out fallibilism is thus to reject EI. As discussed, Certainty infallibilism, where certainty is understood as epistemic certainty, is a version of Ruling Out infallibilism. On the flipside, certainty fallibilism so understood is a version of ruling out fallibilism, and so is likewise incompatible with EI.

Entailment Fallibilism is also incompatible with these principles. The entailment fallibilist holds that S can know that P, even if S's evidence does not entail that P. That is to say, the entailment fallibilist holds that S can know that P, even if some $\sim P$ possibility is compatible with S's evidence. This is in tension with EI, according to which if S knows that P, then $\sim P$ is not possible for S given S's evidence. Dodd (2011) argues that Probability 1 infallibilism is also straightforwardly incompatible with the two epistemic possibility principles. Here is a plausible assumption about probability: if there is a non-zero probability that P, then P is possible. This principle holds across domains (with both the probability and possibility being in the same domain): if the laws of physics assign a non-zero probability relative to physical probability space to the physical possibility that my cup of coffee will spontaneously combust, then it is physically possible that my cup of coffee will spontaneously combust. And if one's epistemic situation assigns a non-zero probability to the possibility relative to epistemic space that Belle is in her office, then it is

¹⁴ See DeRose (1991, 1998), Reed (2013) and Anderson (2014).

epistemically possible that Belle is in her office. If the probability for S that P is non-zero, then P is epistemically possible for S. Contraposing this principle gives us: if $\sim P$ is not epistemically possible for S, then the probability for S of $\sim P$ is zero – and therefore, the probability for S of P is 1. This assumption, when combined with EI, yields infallibilism. For EI says that if S knows that P, then $\sim P$ is not epistemically possible for S. And the assumption now says that if $\sim P$ is not epistemically possible for S, then the probability of P for S is 1. It follows by *modus ponens* that if S knows that P, then the probability of P for S is 1.

Infallibilism in all four forms, then, can explain the infelicity of CKAs: provided that the operators in CKAs quantify over epistemic space then CKAs are infelicitous for the simple reason that they are false. If I know that Alana committed the murder, then it is not epistemically possible that she didn't. And if it's epistemically possible that Belle is in her office, then I do not know that Belle is in the Garden. Infallibilism coheres with the widespread understanding of epistemic modals captured by EI and EP, and as a result preserves the everyday sense in which epistemic modals serve the function of describing what is epistemically possible from the agent's epistemic perspective.

1.3.2. Knowledge does not admit of degrees

One fallibilist solution is to abandon the traditional absolutist view of knowledge and to instead understand both knowledge and epistemic possibility so that they admit of degrees. Baron Reed (2013) defends such a fallibilist account. On Reed's view, epistemic possibility is governed by epistemic certainty, rather than (any degree of) knowledge. If P falls short of being certain for S, P is an epistemic possibility for S. All fallibilist accounts allow that one can clear the threshold for knowing a proposition while falling short of certainty. Reed takes this to suggest that there are many knowledge relations between the bare minimum needed for knowledge and the maximum needed for certainty. Here is how Reed (2013:52, original italics) puts the point:

Once we recognize that we can talk about knowledge without using 'knows', it becomes clear how many different ways there are of doing this. In addition to being *pretty sure* of something, we can also be *very sure* or *highly confident*. We can take something as having been established *beyond a reasonable doubt* or *beyond the shadow of a doubt*. These different comparative terms allow us to refer to our epistemic positions in ways that are more fine-grained than simple knowledge attributions. In some cases, moreover, we do so by setting both upper and lower bounds to the strength of our justification: "I'm not *certain* about it, but I *am* pretty sure.

For Reed, this explains why S can have fallible knowledge that P while $\sim P$ remains epistemically possible for S: $\sim P$ is epistemically possible according to a higher degree of knowledge than the one that S currently bears in relation to P. Furthermore, Reed argues that admitting of degrees is not unique to knowledge. There are many shades of blue one of which is teal. Ordinarily, if x is teal then x is also blue. But there are some cases, for example when choosing a specific team's sports jersey, where one may felicitously utter "that's not blue – it's teal!", where the speaker's intention fixes 'blue' as 'standard blue'. Reed argues that just as 'it's blue' can be false of a jersey

while ‘it’s teal’ is true, so the more general use of ‘knows’ in a sentence can make the sentence false, while a more specific instance – such as ‘pretty sure’ – would make the sentence true. Reed concedes that ‘knows’ *appears* to capture an absolute relation, but argues that the appearance of absoluteness is a linguistic convention and not a reflection of whether the knowledge relation actually is absolute.

In chapter 6 (§3.3), I suggest that Reed’s view gets something fundamentally right: there are some situations in which knowledge attributions are used to communicate some other epistemic relation (e.g. psychological certainty or being supported by good evidence). However, unlike Reed, I do not think that these relations are part of the literal meaning of ‘knows’, but rather are communicated in virtue of sentences of the form “S knows that P” coming equipped with a non-literal meaning. On my view, knowledge does not come in degrees.

Indeed, there is an important disanalogy between Reed’s leading example of *blue* and *knowledge*: when the speaker utters “it’s not blue – it’s teal!” a charitable interlocutor might understand the speaker as denying that the jersey is specific shade of blue, as opposed to misunderstanding them as denying that the jersey is blue (general term). But once the interlocutor has accepted that the jersey in question is in fact teal, we would interpret them as being incompetent with the word ‘teal’ if they continued inquiring into the question of whether it is also blue (general term). The challenge “ok, it’s teal - but is it blue?” is infelicitous to a competent user of the concept teal. In contrast, if I claim to be pretty sure that the party starts at 9pm, the challenge “Ok, you’re pretty sure – but do you know?” is perfectly felicitous, and would not be interpreted as incompetence with either epistemic term. This suggests that knowledge is not a general term denoting different degrees.

There is a more serious reason to deny that knowledge comes in degrees: understanding knowledge as an absolute term best explains the very CKAs that Reed takes to support the knowledge degrees thesis. Reed argues that the ambiguity of the term ‘knows’ explains CKAs in virtue of the speaker’s intention fixing which degree of knowledge she is referring to in each instance. Sometimes, the speaker uses different meanings of ‘knows’ within a single sentence. This sounds infelicitous because generally, speakers have pragmatic reasons to keep their use of ‘knows’ consistent within a given sentence.

If knowledge comes in degrees, then it is plausible that each degree admits of a degree of epistemic possibility, and that speakers can fix the degree with precision. However, there are good reasons to think that knowledge is an absolute term that picks out a single relation. The infelicity of sentences such as 1 and 2 is often used to show that knowledge terms are absolute:

1. Jane knows that Antwerp is cold better than Drew knows it.
2. David knows that the governor will end up in prison, but he doesn’t absolutely know it.

The infelicity of sentences 1 and 2 demonstrates the incompatibility between ‘knows’ and two rules typical of relative terms. 1 shows that ‘knows’ is incompatible with the rule of comparison and 2 show that knows is incompatible with the rule of qualification. Relative terms such as

‘regret’ admit of comparison and qualification: “She regrets selling her car more than she regrets selling her bike”, “He really regrets waking up late”. On the knowledge degrees account these sentences should sound felicitous; 1 would communicate that Jane knows to a higher degree, and 2 would fix David’s knowledge to a degree that falls short of absolutely knowing. The infelicity of these sentences thus lends support to the thesis that knowledge is an absolute term.

As noted, Reed concedes that ‘knows’ appears to capture an absolute relation, but does not take this superficial appearance to reflect the underlying structure of ‘knows’. He argues that we have other ways of talking about knowledge, and these terms illustrate that knowledge is in fact gradable. For example, sentences containing expressions such as ‘pretty certain’, ‘highly likely’ capture knowledge relations while admitting that these knowledge relations vary in strength. Unlike ‘knows’, these terms conform to the rules for determining whether expressions are gradable in terms of whether they admit of qualifiers or comparatives. Consider:

3. Jane is more confident that Antwerp is cold than Drew is.

4. David is pretty sure that the governor will end up in prison.

Further, we can use different expressions to communicate different degrees of knowledge, as in:

5. Charlotte isn’t certain that Anne lives in Pittsburgh, but she is rather confident that Anne lives in Pittsburgh.

For Reed, these are all examples of knowledge relations under different guises. 3 compares Jane’s knowledge to Drew’s, representing that Jane stands in a higher degree of knowledge relative to the proposition that Antwerp is cold. 4 contains a knowledge relation that falls short of certainty while being safely above the minimum threshold. 5 is an example of the common practice of using two separate knowledge relations to define the lower and upper bounds of our knowledge relation. The speaker is not only conveying that she does in fact know, but is specifying the degree to which she knows.

These terms are clearly gradable, but they are not obviously knowledge relations and so do not demonstrate that knowledge relations come in degrees. For a start, none of the terms cited by Reed are factive, whereas ‘knowledge’ is. Given this, Reed’s examples admit of more plausible explanations in terms of agents expressing their non-factive assessment of their psychological levels of confidence. For example, 3 and 4 are plausibly not expressing a knowledge relation at all, but rather represent the speakers’ assessment of the relative psychological certainty or confidence of each of the agents. In 5, both parts of the sentence serve to communicate the speaker’s degree of confidence, explained in terms of an upper and lower limit.

This interpretation along the lines of alleged gradable knowledge expressions in terms of expressing psychological levels of confidence not only explains why the assertions use non-gradable terms, but is also supported by examples of retraction. Consider:

6. David is pretty sure that the governor will end up in prison, but he doesn’t know that he will.

7. Charlotte isn't certain that Anne lives in Pittsburgh, but she knows that Anne lives in Pittsburgh.

6 is felicitous, but 7 does not sound entirely natural. If knowledge were gradable, then we would expect the opposite. If 'pretty sure' were a knowledge relation, then the second clause would be a retraction of the first, and the sentence as a whole should sound infelicitous. 6 does not sound like a retraction because the two parts of the sentence are referring to different phenomena altogether – the first, to psychological facts and the second to knowledge relations. The reverse is true of 7. If 'not certain' can be used to communicate a relation that falls short of the highest degree of knowledge while nevertheless potentially being a lower degree of knowledge, then 7 should not sound like a contradiction.

Additional evidence to suggest that the first half represents a psychological state while the second captures the knowledge relation is found by making the psychological state explicit:

8. David *feels* pretty sure that the governor will end up in prison, but he doesn't know that he will.

9. Charlotte doesn't *feel* certain that Anne lives in Pittsburgh, but she knows that Anne lives in Pittsburgh.

8 is perfectly felicitous. I do concede that to some ears, 9 might sound not entirely felicitous – I think that this is best explained by the tight relationship between psychological certainty and one's perceived strength of epistemic position. If one takes themselves to be in an excellent epistemic position with regards to the proposition that P, then one will also (typically) take themselves to have good reason to feel psychologically certain that P. But cases where these two relations come apart, perhaps, best illustrate that there is no real tension between not *feeling* certain and knowing. What the sentence calls for is some explanation for why the subject's feelings of certainty and knowledge pull apart. For example, in a context where someone had challenged Charlotte on whether Anne does in fact live in Pittsburgh and as a result Charlotte had expressed not feeling certain, the interjection "Charlotte doesn't *feel* certain that Anne lives in Pittsburgh because she's prone to excessively doubting herself when someone challenges her, but she does know that Anne lives there" seems entirely appropriate. On the other hand, 7 remains infelicitous, even with an added explanation.

I think, then, that making the psychological element of these terms explicit by adding 'feels' coheres also with the non-factivity of the terms cited by Reed. Unlike the knowledge relation, psychological levels of certainty towards the truth of a proposition are not factive. What's more, any level of confidence can be substituted into the first half of the sentence, and the sentence as a whole will remain felicitous. Even certainty, when read as psychological, is felicitous. Consider the football fan who on the eve of the finals exclaims, "I feel certain that we will win!" without thereby committing herself to knowing that her team will win. Thus, examples of retraction and examples in which the psychological aspect have been made explicit to suggest that, contra the view that there are different knowledge terms that communicate different degrees of knowledge, the examples cited communicate psychological states being contrasted with knowledge states.

Returning to CKAs, this view of psychological representation suggests that there are two readings to the example we started with:

10 I know that Alana committed the murder, but it's possible that she didn't

11 I feel confident that Alana committed the murder, but it's possible that she didn't

If both clauses capture the knowing relation, then 10 reads as a straight retraction. The speaker started by representing herself as knowing, and then retracted the claim to know in light of considering epistemic possibilities that had not been ruled out. On the other hand, 11 presents no contradiction. The speaker represents her psychological level of confidence on the one hand, and her knowledge as it governs epistemic possibility on the other hand. In any given conversation, the appropriate reading will most likely depend upon the context and the speaker's intention.

To summarise, interpreting alleged gradable knowledge relations as representing psychological confidence rather than knowledge is a more plausible account of the gradable terms cited, and coheres both with other linguistic data that suggests that knowledge is absolute and with the non-factivity of these terms. As a further point to note, even granting that knowledge comes in degrees would have been but a partial victory for the fallibilist. At each level and for each degree, S's knowledge would have been infallible in the sense that it would have admitted of no possibility of error, as defined to that particular degree. As such, knowledge coming in degrees is perhaps best understood as degrees of infallible knowledge, conferring only a limited advantage to the fallibilist. Nevertheless, there is something interesting and essentially correct about the view that knowledge comes in degrees, which I will return to in chapter 6: 'knowledge' is sometimes used to communicate other epistemic relations.

1.3.3. Knowledge is Qualitatively Different to Non-knowledge

The vast majority of analyses for knowledge take the form of a series of conditions which, when combined with belief and truth, are necessary and sufficient for knowledge. These conditions sometimes take the form of 'justification' or 'warrant'.

Examples include:

- a. S's belief that P is sufficiently probable on S's evidence.
- b. S's belief is produced by a reliable cognitive process.
- c. S can rule out all relevant alternatives to P
- d. Were P false, S would not believe that P
- e. S's true belief that P is not due to epistemic luck
- f. If S were to believe P, P would be true
- g. S's belief that P is apt ¹⁵

¹⁵ I here borrow from, and expand on, Climenhaga's (2021:59) list.

Climenhaga (2021) notes that many of these conditions contain threshold terms: ‘sufficiently probable’, ‘reliable’, ‘relevant alternatives’. The ones that do not strictly speaking contain threshold terms are nevertheless structured in ways that admit of thresholds. For example, although “Were P false, S would not believe that P” seems like an all-or-nothing matter, fleshing out the details will involve applying threshold terms. For example, whether this condition obtains will hinge on whether S does not believe that P in the closest worlds to the actual world (and if so, how close must a world be to count as among the ‘closest’ worlds?), or perhaps additionally on whether S does not believe that P in most of the worlds in which P is false (and if so, how many worlds count as ‘most’?). Likewise, whether S’s true belief is the product of S’s exercising competency will be a matter of degree, with some threshold needed to be cleared in order for “S’s belief that P is apt” to be true.

There seems to be a qualitative difference between knowing that P and having a highly justified true belief that P, a difference which fallibilism struggles to explain. For the fallibilist holds that justification comes in degrees, with one particular degree being necessary for knowledge. But this does not explain why knowledge is qualitatively different to justified true belief. As Climenhaga (2021:59) puts it, “the move from 0.89 to 0.9 probability, or truth in the five nearest possible worlds to truth in the six nearest worlds, and so on, is a quantitative one, and not a qualitative one”. Infallibilism, on the other hand, easily explains this qualitative difference: if infallibilism is true, then the difference between knowledge and even highly justified true belief is that only in the case of knowledge is S’s belief that P maximally justified, where this might mean that S has eliminated all of the possibilities of error, or that S’s evidence entails that P, or that S is epistemically certain that P.

One option is to deny that the difference between highly justified true belief and knowledge is qualitative, as opposed to quantitative. Perhaps ‘knowledge’ is a threshold term which does not in fact reflect underlying differences in quality. Other examples include *century* and *adulthood*. There is no substantial or significant difference in quality between 99 years and 100 years, or between being 17 and being 18 years of age underpinning the convention of referring to only one of these as a ‘century’ or an ‘adult’. Rather the threshold marks some culturally or socially significant quantitative difference, such that there has been a need for a concept to demarcate a particular quantitative difference.

However, there does seem to be a qualitative difference between knowledge and highly justified belief. Some theorists argue that there is an intuitive connection between knowledge and rational action, so that it is rational to act on what one knows¹⁶. For example, Fantl and McGrath (2002:72) write that “If you know that P, then it shouldn’t be a problem to act as if P. If it is a problem to act as if P, you can explain why by saying that you don’t know that P”; Hawthorne (2004:30) writes that “it is acceptable to use the premise that P in one’s deliberations if one knows it and (at least in very many cases) unacceptable to use the premise that P in one’s practical reasoning if one doesn’t know it”; and Stanley (2005:10) writes that “A standard use of knowledge attributions is to justify action [...] To say that an action is only based on a belief is to criticize that

¹⁶ See for example Fantl & McGrath (2002), Hawthorne (2004), Stanley (2005), and Hawthorne and Stanley (2008).

action for not living up to an expected norm; to say that an action is based on knowledge is to declare that the action has met the expected norm.” If justified (true) belief is insufficient for action, but knowledge is always sufficient for action, then this suggests that there *is* a difference in quality and not merely quantity between justified belief that P and knowledge that P, with only the latter being of sufficient quality to always be used as a premise in one’s practical deliberations.

Furthermore, BonJour (2010) argues for another key qualitative difference between knowledge and justified true belief: knowledge is valued more than justified true belief. BonJour (2010:58) characterises knowledge as a “supremely valuable and desirable cognitive state, one whose possession marks the difference between full cognitive success and at least some degree of cognitive failure: knowledge is the epistemic *summum bonum*”. Since the difference between justified true belief and knowledge is, for the fallibilist, a matter of degree – a matter of *quantity* – the fallibilist cannot explain what gives “the supposed “magic” level of justification the kind of special significance that it is supposed by [fallibilism] to have” (BonJour 2010:62). Put differently, the fallibilist cannot explain why the difference between justified true belief and knowledge is so much greater than the difference between the various degrees of justification; it cannot explain why we care so very much about acquiring the particular degree of justification that is demarcated by ‘knowledge’. In contrast, the infallibilist is perfectly positioned to explain the difference: the difference in value is a natural consequence of the difference in quality between justified true belief and knowledge. Specifically, knowledge is valuable because it demarcates the level of justification which guarantees truth, and at which the justification could not support a false belief.

1.3.4. Knowledge is Closed under Known Entailment

The principle of closure in its various iterations has enjoyed popularity among epistemologists. To say that knowledge is closed under known entailment, for example, is to say that S is in a position to know every proposition that S knows to be entailed by another proposition that S knows. Put formally, closure says that if S knows that P, and S knows that P entails Q, then S knows that Q. Closure reflect a typical way in which people in everyday contexts extend their knowledge by using deduction. The following ordinary piece of reasoning, for example, employs the principle of closure: I know that Belle is in her office; I know that if Belle is still in the office, she won’t be home in time for dinner; I know that Belle won’t be home in time for dinner. It seems that in everyday life we typically employ this principle to add to our stock of known propositions. We build on known propositions (that Belle is still in the office) via deduction to generate new knowledge (that Belle won’t be home in time for dinner).

It has been widely noted that the principle of closure is incompatible with fallibilism (BonJour 2010, Reed 2012, Climenhaga 2021), and so fallibilism comes at the cost of denying this highly intuitive principle. For if fallibilism is true, then it is possible for S to know P, deduce Q from P, and yet not know that Q. To use Dretske’s (1970) famous zebra case, if I know that the animal in a pen is a zebra and I know that if the animal in the pen is not a cleverly disguised mule then, by closure, I should be able to extent my knowledge and come to know that the animal in the pen is not a cleverly disguised mule. But many epistemologists have thought that I cannot know that the animal in the pen is not a cleverly disguised mule. Fallibilism, then, comes at a cost: either the fallibilist must deny the highly intuitive principle of closure, or she must say that it is

possible to know that the animal in the pen is not a cleverly disguised mule. Infallibilism does not face this problem, since on infallibilism if S knows that P and deduces that Q is entailed by P, then one also (infallibly) knows that Q¹⁷.

1.4. Conclusion

This chapter has tried to motivate the claim that there are good reasons to think that sceptical infallibilism about knowledge is correct. Not only does sceptical infallibilism capture the intuitive sense in which if S knows that P, then S cannot be mistaken about P, but infallibilism is explanatorily powerful: it is compatible with plausible theories on epistemic probability, and hence explains otherwise puzzling concessive knowledge attributions; it explains why there seems to be a qualitative difference between knowledge and justified true belief; and it explains why the principle of closure is intuitive. Further, I have argued that however we conceive of infallibilism, at least on one plausible iteration, infallibilism generates high-standards scepticism. Both friends and foes of infallibilism treat the sceptical consequences of infallibilism as a significant cost. This thesis develops and defends a hitherto underexplored view: epistemic fictionalism, according to which sceptical infallibilism is (at least on one level) correct, but on which the negative consequences of scepticism are avoided. Before turning to epistemic fictionalism, Chapter 2 addresses a relatively new objection to sceptical infallibilism: that epistemic infallibilism does not cohere with the purpose of the concept of knowledge. This chapter has the additional aim of introducing the methodology central to this thesis: function-first epistemology.

¹⁷ Multi-premise closure principles also pose a problem for fallibilists, as illustrated by preface paradoxes: according to fallibilism, S can know each of the 1000 claims in her book, even if on S's evidence the probability of each being true is less than 1 – say that the probability of each being true is 0.99. Applying the principle of closure, S should be in a position to come to know the proposition “all of the claims in this book are correct”. But the probability of all of the claims in S's book being true is extremely low (0.0004, assuming that the claims are independent of one another). Fallibilism bears the cost of either denying multi-premise closure, or having to claim that one knows that all the claims in the book are correct. As was the case for single-premise closure, infallibilism does not face this problem: if S knows each of the 1000 claims in her book then the probability of each being true is 1, and the probability of “all of the claims in this book are correct” is also 1.

Chapter 2: Function-First Epistemology and the New Apraxia Objection to Sceptical Infallibilism

Several epistemologists have recently defended versions of the New Apraxia objection to sceptical infallibilism, according to which sceptical infallibilism is false because it is incompatible with the point of knowledge (Fricker 2008, Kaplan 2008, Henderson 2011, Hannon 2019b and Queloz 2021). This chapter outlines the New Apraxia objection (§2.2.), before developing a response to this objection on behalf of the sceptical infallibilist. I argue that positing systematic and useful error makes the sceptical infallibilist's concept of knowledge just as well-equipped to fulfil knowledge's functions as the anti-sceptic's concept of knowledge (§2.3.). To illustrate this point, I discuss two errors that sceptics can (and typically do) appeal to, and show how positing either one of these two errors enables sceptical infallibilism to meet the New Apraxia objection. Since the sceptical infallibilist can account for how knowledge fulfils its functions, the New Apraxia objection fails. This chapter also serves the purpose of introducing the function-first method that underpins this thesis (§2.1.).

2.1. Function-First Epistemology

In epistemology, the view that investigating the function of the concept of knowledge can illuminate our understanding of knowledge is gaining traction¹⁸. An epistemologist who adopts the function-first method seeks to understand or evaluate a concept by exploring what the point of the concept is: what having this concept as part of our conceptual repertoire enables us to do that we could not otherwise do (or could not otherwise do with as much ease, or as efficiently). The thought is that all of our concepts, knowledge included, emerge and remain in circulation because they meet particular needs or enable us to achieve something that we could not, or could not so easily, achieve without that concept. The function-first philosopher is concerned with what this 'something' is: by what needs, purposes or functions are met by the target concept. If we are interested in investigating a concept's intension and extension, we may proceed by first identifying what function the target concept fulfils, and then ask what intension and extension the concept would require in order to successfully fulfil the posited function.

The term 'Function-First Epistemology' owes its origins to Michael Hannon (2019a), who describes it as a methodology which "seeks to explain the nature and value of an epistemic concept, norm, or practice by reflecting on its function or purposes" (2019a:12). Thus a function-first epistemologist asks questions like (2019a:12):

Why do humans speak and think in terms of "knowing," "understanding," and "rationality"? What epistemological distinctions and norms would best facilitate human survival and flourishing? What might life be like without our current

¹⁸ See for example Craig (1990), Hannon (2019a, 2019b), Kvanvig (2003), Fricker (2008), Kusch (2009), Henderson (2009), Kelp (2011), Rysiew (2012) and Beebe (2012) for epistemologists who centre their epistemology around the point of knowledge.

practices of epistemic evaluation? Do our concepts of epistemic evaluation carry much weight in science, philosophy, or daily life?

Function-first epistemology goes under a variety of different names, including ‘purposeful epistemology’ (Henderson and Greco 2015:1) and ‘teleological epistemology’ (Gardiner 2015:31). Edward Craig, whom many regard as the first epistemologist to develop and apply the Function-First methodology to the concept of knowledge¹⁹, calls his method ‘practical explication’ (Craig, 1987, 1990), and describes it as follows (1987:2012):

Instead of beginning with ordinary usage, we begin with an ordinary situation. We take some *prima facie* plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like.

Function-first epistemology is concerned with the concept of knowledge. For my purposes it does not matter whether there is also a state of knowledge that may be distinguished from the concept: what might be thought of as “knowledge as such” (Hannon 2019a:32) or “knowledge itself” (Queloz 2021:132). If there is such a state, then it is still insightful to investigate why this particular state has been especially picked out with the concept of knowledge. Furthermore, I do not wish to commit to any particular view of the nature of concepts, other than to note that concepts are constituents of thought and talk. In order for one to think “there is a cat on the mat”, one must have the concepts of *cat* and *mat*. Furthermore, I take concepts as they appear in fragments of discourse to have the same content (for example, the same intension and extension) as they have in thought. Thus an utterance of “there is a cat on the mat” expresses the same content as the thought “there is a cat on the mat”. This thesis concerns the concept of knowledge, especially as it appears in fragments of discourse (e.g. in sentences of the form “S knows that P”), which I refer to as ‘knowledge talk’. Examining the function of the concept of knowledge, then, is to also examine the function of knowledge talk. For as McGrath (2015:138) notes, “[w]hat a concept does for us is a matter of what we can do with it, and what we do with concepts is deploy them in thought and language”.

It is perhaps most helpful to note what function-first epistemology is *not*, by way of comparing it to two more familiar philosophical methods: the *intension-first* and the *extension-first* approach to concepts (see Gardiner 2015: 32; Queloz 2021: 44). A philosopher engaged in an intension-first analysis assumes that we have intuitive access to the intension of our concepts, and uses this intuitive intension both as a starting point in her investigation and as the central datum for the finished analysis to capture. For example, a philosopher who applies an intension-first approach to investigate knowledge might aim to construct necessary and sufficient conditions by first gathering both folk and theory-driven platitudes about knowledge’s intension: that knowledge implies truth, that knowledge involves belief, that knowledge requires strong evidence, and so on. Whichever platitudes seem most central to knowledge, she will build into necessary and sufficient conditions for satisfying the concept, turning her intuitive intension into an explicit

¹⁹ See Hannon (2019a), Fricker (2008), Henderson and Greco (2015), Gardiner (2015), Kusch and McKenna (2020).

intension: for example, knowledge is true belief based on strong evidence. She will then use this intension to find knowledge's extension: she will look to the world for instances of true belief based on strong evidence, and class all instances as falling under the extension of knowledge. Chapter 1 predominantly applied this methodology, arguing that folk and theoretical platitudes about knowledge's intension – such as the infelicity of CKAs and the idea that knowledge is qualitatively different from non-knowledge – support infallibilism.

A philosopher engaged in an extension-first analysis assumes that we have intuitive access to the extensions of our concepts: that we can intuitively tell what objects fall under the concept. She will centre her project on an exploration of its intuitive extension, and test proposed intensions for the concept by how well they capture this intuitive extension. Counterexample-trading is a key method of the extension-first methodology. Consider the post-Gettier literature on knowledge, where increasingly convoluted examples to proposed intensions of knowledge are offered: Gettier's (1963) original coin and Ford cases that constitute counterexamples to the intension of knowledge as justified true belief; the fake barns counterexample (Goldman 1976 credits Carl Ginet with coming up with the example) to Alvin Goldman's (1967) causal account of knowledge; Ernest Sosa's (1999) trash chute counterexample to Fred Dretske (1971) and Robert Nozick's (1981) sensitivity conditions; and so on. In each case, the presence of some instance of belief that satisfies the necessary and sufficient conditions for knowledge contained in the proposed intension for the concept, but which intuitively does not fall under the extension of the concept, is taken as reason to modify or abandon that proposed intension²⁰. Fallibilists who seek to resist infallibilism on the basis that it engenders scepticism are typically applying extension-first methods, arguing from the intuition that we know a lot – i.e. that many objects fall under the extension of the concept of knowledge – to the conclusion that sceptical infallibilism is false.

In contrast to both of these approaches, a function-first philosopher need not assume that we have intuitive access to the intension or extension of our concepts, and does not centre her investigation into the target concept on intuitive intension or extension. Rather, the function-first philosopher starts by applying philosophical methods to work out what function a concept plays. The function-first philosopher might start her investigation by simply stipulating a *prima facie* plausible hypothesis regarding the function of her target concept, and then test it by seeing how well a concept constructed to meet those needs would match our current concept (Hannon 2019b). Or she might employ the method of conceptual reverse-engineering, which aims to find the function of a given concept by reconstructing the practical problems that some community has to which the concept offers a solution (Queloz 2021:53)²¹.

²⁰ The characterisation of intension- and extension- first approaches, and supporting examples, are drawn from Newton and O'Sullivan (2023, manuscript).

²¹ This was of understanding conceptual reverse engineering as a tool for function-first epistemology is somewhat in tension with Hannon's (2019a) way of viewing the landscape, on which conceptual-reverse engineering is a different (although perhaps complimentary) methodology. For Hannon, the key difference between the two methods is the direction of investigation: function-first starts with function, to investigate the concept; conceptual-reverse engineering starts with the concept, in order to investigate the function.

The philosopher engaged in conceptual reverse-engineering has many and various tools at her disposal for reconstructing the problem to which the target concept offers a solution. For example, she might use a model, either synchronic (Queloz 20221 offers Wittgenstein 2009 as an example) or diachronic. Diachronic models might take the form of a genealogy, which in turn might be imaginary, as in state-of-nature stories (e.g. Hume 1739, Craig 1990 and Williams 2002) or historical (e.g. Foucault 1971, Gelfert 2011)²². Not all conceptual reverse-engineers are modellers. Some use what, adopting Weisberg's (2007) terminology, we can think of as abstract direct representation: abstracting away from the particular features of real scenarios to represent the core needs that the concept emerges to meet. These might also be synchronic, taking the form of what Fricker (2016) calls 'paradigm-based explication' on which one focuses on an abstracted version of a typical case in which the need is present; or they might be diachronic, abstracting from real cases across a period of historical time (e.g. Nietzsche on punishment)²³. Irrespective of how the function-first philosopher generates her initial hypothesis regarding the function of the concept, it is this hypothesis that drives her investigation into the concept, rather than intuitions regarding intension and extension.

Once she has settled on a hypothesis, the function-first philosopher works backwards to consider what a concept would need to look like in order to successfully fulfil this function. She might consider what conditions govern the application of the function; what intension and extension it might have; how the concept interacts with other related concepts; and whether there are any puzzles surrounding the concept where progress might be made by considering whether there are solutions that block, or promote, the fulfilling of the function. One important upshot is that a function-first analysis need not take the form of necessary and sufficient conditions, but rather picks out how the concept has developed to function in typical cases. A concept develops to meet particular needs, in a particular environment. Compatibly with this, there may well be cases so dissimilar from the typical environment in which the concept developed that it is indeterminate whether the concept applies. For example, Craig (1990: 52) argues that Gettier cases are such examples: they are so unusual, both in the sense of taking place in an untypical environment and in the sense of being reflected on from an untypical third-person perspective, that our concept of knowledge has not been shaped to stretch over these cases. The main point is that so long as they are untypical, counterexamples does not put pressure on a function-first analysis. In short,

²² See Queloz (2021) for a thorough defence of the idea that genealogies are fruitfully conceived of as models

²³ Understanding Nietzsche's account of punishment as a diachronic abstract direct representation is in my view the most charitable reading of Nietzsche's project (Newton and O'Sullivan 2023, manuscript). Queloz (2021) notes that although Nietzsche is typically read as giving a real history of punishment, his lack of real historical references suggests that we should instead understand him as offering an imaginary state-of-nature account; to this, Hannon (2022) objects that Queloz's analysis cannot make sense of Nietzsche's consistent use of causal-explanatory language. But thinking of Nietzsche as offering an *abstraction* of real historical events, analogous to Fricker's abstraction of present day cases, explains both of these features. Nietzsche abstracts away from the particular features of real historical events to describe their core features. In doing so, he is nevertheless describing real features of real history— it is just real history stripped of details that are irrelevant to the core function of the concept under investigation. He is not a modeller, because he is describing real historical events (albeit, abstracted); but neither is he engaged in historiography, because he is not faithfully recounting events at a particular place and time, but abstracting from the core features of multiple real events.

the function-first epistemologist is “interested in identifying what might be called the “core” of our epistemic concepts and practices”, where this might take the form of “conditions that hold only for the most part, but not always” (Hannon 2019a:19).

Nevertheless, a focus on function does not mean that the function-first philosopher is free to disregard intuitive intension and extension altogether. On the contrary, the way in which a function-first philosopher verifies her analysis is by testing it against intuitive intension and intuitive extension. For example, Craig (1990:2) writes the following regarding function-first investigations:

Such an investigation would still have an anchorage point in the everyday concept: should it reach a result quite different from the intuitive intension, or one that yielded an extension quite different from the intuitive extension, then, barring some special and especially plausible explanation of this mismatch, the original hypothesis about the role that the concept plays in our life would of course be the first casualty.

Hannon (2019a:14) echoes this sentiment, writing that:

Our aim is to construct a concept that not only functions in the manner suggested by our hypothesis, but also one that fits our intuitions (or explains why our intuitions diverge). Thus, while function comes “first” on this approach, it is not the last word.

So intuitive intension and intuitive extension play an important, if secondary, role. Spelling out exactly what ‘first’ and ‘second’ mean is not entirely straightforward. Hannon typically writes as if the sense in which function comes ‘first’ is literal, representing the first stage in a structured investigation; and sometimes he writes as if ‘first’ is metaphorical for explanatory priority. I think the second sense in which function comes first is the more illuminating. Putting function first means that we understand intension and extension as *resulting* from the concept serving a function. The function is the thing that explains the concept having the intension and extension that it has. Where there is tension between one’s function-first analysis of a target concept on the one hand, and its intuitive intension and extension on the other, then the analysis is successful to the extent to which the function explains why the intuitive intension and extension are as they are. This is what Craig means when he asks for a plausible story for how the “massive misfit could have come about” (1990:2): an explanation, ground in function, for why a concept might acquire an intension or extension that differs significantly from our folk conception of either intension or extension. In a similar vein, where intuitive intension and intuitive extension pull apart, or where there are a range of intuitions among everyday concept users and philosophers about what the intuitive intension and extension is, the function-first philosopher tries to put function to work in diagnosing the source of this divergence (Craig 1990, Hannon 2019a, Kusch and McKenna 2020). To do this, the function-first philosopher might ask: why has the concept emerged so as to generate these competing intuitions? How does this ambiguity enable to concept to fulfil its functions?

Finally, this sense in which function comes first is also reflected in the function-first philosopher’s commitments, measured in terms of which aspects of the concept the philosopher is willing to

surrender when there is tension between function, intuitive tension and intuitive extension. Whereas philosophers engaged in intension-first or extension-first philosophy automatically reject analyses which stand in tension with the intuitive intension and extension respectively, the function-first philosopher automatically rejects analyses which stand in tension with the posited function. In this way, the hypothesis about the function guides the analysis, and the function-first label denotes “the relative importance of functional considerations in our theorizing” (Hannon 2019a:26).

2.2. Anti-sceptical Function-First Epistemology

With this methodology on the table, let’s return to the question of sceptical infallibilism. One of the advantages that function-first epistemologists claim for their method is its alleged success in refuting the sceptical infallibilist. In a nutshell, the idea is that the sceptical infallibilist’s concept of knowledge comes equipped with such demanding standards that it always (or almost always) fails to apply. Such a concept would be pointless, for no practical purposes call for a concept that fails to apply in all (or almost all) of the circumstances in which speakers attempt to apply it. For this reason, several function-first philosophers have argued that taking seriously the idea that knowledge emerged to fulfil a function undermines sceptical infallibilism. If sceptical infallibilism is incompatible with any plausible hypothesis regarding the function of knowledge then, as per the methodology outlined above, we should reject sceptical infallibilism.

2.2.1. *The New Apraxia Objection to Scepticism*

Hannon (2019b:144) terms this anti-scepticism strategy the ‘New Apraxia’ objection to scepticism. New Apraxia contrasts with traditional apraxia objections to scepticism, on which scepticism is to be resisted because it would be practically disadvantageous to live as a sceptic. Instead, the New Apraxia objection holds that the truth of scepticism is incompatible with knowledge serving a function at all. Since by assumption knowledge does serve a function, sceptical infallibilism is thereby refuted.

We find this idea right at the origins of function-first epistemology, in Edward Craig’s (1990) *Knowledge and the State of Nature*. For Craig, knowledge serves the function of recommending reliable informants (see also Hannon 2019a). Craig reasons that in order to survive, each individual needs to gather information regarding their surroundings, and sometimes other members of the community offer the best source of information. However, on any topic some members of the community are more likely than others to be in a position to provide reliable information. The concepts of knowledge and of knower emerge to enable us to ‘flag’ these members of the community in order to identify them and recommend them to one another. Craig notes how in most everyday contexts, the sceptic’s standards for knowledge are so demanding that even highly reliable informants would fail to qualify as knowers, thus making knowledge ill-suited for the everyday purpose of recommending informants. Indeed, would-be recommenders employing the sceptic’s knowledge concept would withhold knowledge attributions from reliable informants, misleadingly conveying to their audience that said informants are poor informants. Craig grants that perhaps being engaged in certain specific and unusual inquiries, for example, Descartes’ search for certain foundations, pushes the standards

for knowledge upwards towards those posited by the sceptic. Nevertheless, the sceptic who takes these maximally demanding standards to be operative in everyday contexts is mistaken, since everyday contexts call for standards no more demanding than those met by reliable everyday informants.

Several other function-first epistemologists follow Craig in extracting an anti-scepticism lesson from hypotheses about the function of knowledge. Henderson argues that evaluative concepts such as knowledge emerge to fulfil the function of regulating the behaviour of agents intent on pursuing specific projects, and therefore what is required to satisfy evaluative concepts is constrained by what is required to achieve the aims of the project. It follows that if “one’s conception of knowledge would frustrate the individual and community practices that the concept of knowledge grew up to regulate, one’s conception would need to give way” (2011:92)²⁴. Fricker argues that the practical everyday context of recommending informants forms the ‘core’ of the concept of knowledge, so that even sceptics “must be inquirers first and sceptics second” (2008:48). That is to say, the cases that best demonstrate the functionality of knowledge are cases in which the standards for knowledge are attainable. Hannon (2019b) defends at length the claim that “scepticism is implausible because it goes against the point of epistemic evaluation” (2019b:154). He argues that knowledge bearing an empty extension would prevent knowledge from fulfilling any of the four most popular candidates for the function of knowledge (recommending informants, ending inquiry, regulating assertion and regulating practical reasoning), and so scepticism is to be rejected. In a similar vein, Kaplan argues that the sceptic’s knowledge concept is misguided due to being “hopelessly inadequate” (2008:358) for enabling concept users to enjoy an established policy on how to act when the truth of a proposition has practical import for a decision, which in turn enables them to share useful information and expertise. Finally, Queloz writes that the sceptic’s knowledge concept “would be utterly pointless, since it would simply fail to get a grip on the world we inhabit—nothing would ever satisfy the concept” (2021:150). Queloz goes on to argue that knowledge cannot have been equipped with the standards posited by the sceptic when knowledge first emerged. If the present knowledge concept is as the sceptic posits, it must have lost its original function along the way.

Many and various functions have been posited for the concept of knowledge. In addition to the ‘reliable informant’ function endorsed by Craig (1990), Hannon (2019a) and Pritchard (2012), epistemologists have argued that knowledge functions to signal that inquiry may legitimately come to an end (Kappel 2010, Kelp 2011, Rysiew 2012); that inquiry into a question may begin (Fricker 2008); of distinguishing between blameless and blameworthy behaviour (Beebe 2012); of reaching collective agreement for joint decision-making (Gelfert 2011); of honouring information gatherers (Kusch 2009); of identifying propositions that can be treated as reasons for acting (McGrath 2015); of providing assurance (Austin 1946, Lawlor 2013)²⁵. The New Apraxia objection against the sceptical infallibilist appears to hold irrespective of which function we posit for knowledge. Take the inquiry-stopping function. The idea behind this hypothesis is that

²⁴ Henderson (2011) follows through by revising his own (2009) function-first analysis of knowledge in light of its scepticism-friendly results.

²⁵ Many function-first philosophers (Beebe 2012, Kusch 2009, McGrath 2015, Hannon 2019a, Gelfert 2011) posit multiple, complementary functions.

inquiry into any question is potentially infinite: given unlimited time and resources, we could keep on searching for better and better evidence indefinitely. However, creatures like us with limited time and resources require a point at which to legitimately cease inquiring. Knowledge serves the function of signalling the point at which, given the context and the costs of further inquiry, one may stop inquiring. But on sceptical infallibilism, the New Apraxia objection goes, no one is ever in a position where they may stop inquiring and therefore the sceptical infallibilist concept of knowledge would not be able to fulfil its function of stopping inquiry.

Hannon frames the New Apraxia objection in terms of a dilemma for the skeptic: “[e]ither the purpose of knowledge is such that it rules out skepticism, or else knowledge is itself a useless and uninteresting concept” (2019b:147). He argues that the sceptic’s knowledge concept is useless because incapable of fulfilling a function, and uninteresting because it would fail to describe our actual knowledge concept. Perhaps the sceptic can stretch the concept of knowledge, or create her own concept anew, but the discovery that this new concept always or almost always fails to apply should not be alarming to us. What we are interested in investigating is whether our actual knowledge concept applies. Because we have reason to think that our knowledge concept serves a function, the argument goes, we have reason to think that it does apply. Thus Hannon urges us to embrace the first horn of the dilemma: anti-scepticism.

Insofar as one wishes to remain in the function-first programme and describe our actual knowledge concept, one cannot accept the second horn: that knowledge is useless and uninteresting. I do not, however, think that we are forced to accept the first horn of the dilemma, either. There is a third way out: argue that the function of knowledge is compatible with the truth of scepticism after all. For sceptics taking this route, Hannon’s dilemma can more fruitfully be reformulated as a constraint on sceptical theories: the sceptic’s knowledge concept must be both useful and interesting. It must be useful in that it must be capable of fulfilling a function, and it must be interesting in that it must be a description of our actual knowledge concept. Not only is this a plausible and welcome constraint on sceptical theories but sceptics can, and indeed typically do, meet this challenge. Or so I argue.

2.2.2. Double Bookkeeping

Here is a *prima facie* intuitive way in which the sceptic might explain how her knowledge concept successfully fulfils a function: divorce the attribution conditions for knowledge from the truth conditions for knowledge. After all, sceptics tend to agree with anti-sceptics that it is often appropriate and pragmatic to attribute knowledge - the dispute between sceptics and anti-sceptics concerns whether these knowledge attributions are typically true, not whether they are typically useful. Butchvarov (1970), Unger (1975), Stroud (1984), Stone (2000), Schaffer (2004), Dodd (2011), Bonjour (2010), Russell (2021) and Stoutenburg (2021) all supplement their sceptical arguments with a pragmatic account of why it is expedient to attribute knowledge in everyday life, despite these knowledge attributions being typically false. Say that knowledge is useful because it enables us to recommend reliable informants. Grant that the sceptical infallibilist is correct: knowledge comes equipped with maximally demanding standards and so knowledge attributions are typically false. Yet, so long as the conditions for attributing knowledge reliably and consistently cause knowledge to be attributed to sufficiently reliable informants, the concept of knowledge

can successfully fulfil its function. Put differently, sceptical infallibilists might seek to resist the New Apraxia objection by conceding that fulfilling the function of knowledge requires that the conditions for attributing knowledge are those predicted by the anti-sceptic, while denying that the truth conditions for knowledge need to likewise be anti-sceptical.

Kyriacou's (2020) 'pragmatic scepticism' is an explicit example of this strategy. Kyriacou argues that since unattainable norms of assertion and practical reasoning would have left our "everyday practice of assertion and practical reasoning all but ruined" and our "practical and intellectual lives in tatters" (2020:554), practical pressures push the attribution conditions for knowledge towards achievable, anti-sceptical standards. Compatibly with this, however, "theoretical reasons" push the truth conditions towards those posited by the sceptic. Because the systematically false knowledge attributions nevertheless work "reasonably well for all practical purposes" (2020:549), knowledge successfully fulfils its practical roles. Kyriacou's pragmatic scepticism appears to preserve the interest and utility of knowledge: it elucidates our knowledge concept, and explains how knowledge fulfils a function.

Hannon is pessimistic about this style of response to the New Apraxia objection. He objects that investigating knowledge from within the function-first programme means commitment to the claim that both the truth conditions and attribution conditions for knowledge are shaped by the practical needs which knowledge has developed to meet. If knowledge functions to meet a practical need, then why think, as Kyriacou does, that entirely different theoretical considerations shape its truth conditions? Hannon argues that there are only two explanations for why truth conditions and attribution conditions might come apart, neither of which a function-firster should accept: either one single set of practical needs shape the truth conditions and attribution conditions in opposite directions, saddling the sceptic with the onerous task of explaining why; or practical needs do not shape the truth conditions at all and the sceptic is forced to explain why knowledge would come equipped with truth conditions that are completely divorced from the point of the concept. As Hannon (2019b:152) puts the point:

Skepticism leaves unexplained why the assertability conditions for "know" would be sensitive to practical exigencies while the truth conditions for knowledge would be isolated from, and unaffected by, such pragmatic considerations. Why would the conditions for knowing be isolated from the utility of this practice? In short, why keep two sets of books?

Call this the 'double bookkeeping objection'. The idea is that it is implausible to suppose that practical needs either fail to shape, or differently shape, the truth conditions for knowledge from the attribution conditions for knowledge. As Hannon emphasises, the double bookkeeping objection does not make scepticism logically incoherent. Nevertheless, it does add a further constraint to a response to the New Apraxia objection and, by extension, to a successful function-first sceptical infallibilist analysis of knowledge: the sceptical infallibilist must either defend the claim that the truth conditions for knowledge are untouched by the functions which knowledge fulfils, or must provide a plausible explanation for why the same practical needs shape the truth conditions and attribution conditions differently. I agree with Hannon that the first option is incompatible with a function-first analysis of knowledge; but the sceptical infallibilist can offer a

principled explanation for double bookkeeping. Indeed, I argue that there are two distinct ways in which to make double bookkeeping plausible on a function-first analysis of knowledge, each already invoked by at least one prominent scepticism-friendly view.

2.3. Useful Error

How, then, do practical (as opposed to theoretical) pressures coexist with knowledge coming equipped with the truth conditions posited by the sceptical infallibilist and with different, anti-sceptical attribution conditions? The sceptical infallibilist can answer this question by appealing to useful error, that is, to particular patterns of error that yield advantages to communities that systematically make them. In a nutshell, the sceptical infallibilist argues that practical pressures are compatible with the truth conditions posited by sceptical infallibilism, provided that systematic patterns of error prevent speakers and hearers from realising that the truth conditions of their knowledge attributions typically go unmet. Indeed, every sceptical theory posits systematic error of one type or another. Here are two errors which are popular with sceptics: either we are mistaken about the strength of our epistemic position, erroneously taking our beliefs to meet knowledge's maximally demanding standards and therefore erroneously taking the maximally demanding truth conditions for knowledge attributions to frequently be met; or we are mistaken about the truth conditions for "S knows that P", erroneously taking them to be less demanding than they are.

The truth conditions for knowledge attributions posited by the sceptic, combined either with systematic error regarding whether the truth conditions posited by the sceptic are met or with systematic error regarding how demanding the truth conditions are, yield the same pattern of practically useful knowledge attributions as would error-free non-sceptical truth conditions. To illustrate how the function of knowledge is compatible with double bookkeeping, this section examines a version of each of these types of error, and argues that they successfully explain why knowledge's truth conditions and attribution conditions come apart. The aim is not (yet) to endorse either one of these views, but rather to illustrate two ways in which the sceptic might meet the double bookkeeping challenge and respond to the New Apraxia objection.

Before discussing these two views, I wish to clarify the role that error plays in the defence of scepticism that I advance. Sceptical analyses of knowledge typically begin by noting the theoretical advantages of scepticism, and later posit error in an attempt to explain away the mismatch between the theoretically supported sceptic's theory of knowledge and observable everyday patterns of knowledge attribution. The underlying assumption is that scepticism is correct on philosophical or theoretical grounds, but that positing widespread error is somehow uncharitable to knowledge concept users. So, after advancing the theoretical arguments for scepticism, the sceptic must explain away the posited error that she attributes to everyday speakers. For example, Stroud (1984) argues that, among other theoretical virtues, scepticism respects the principle of closure; yet people do not realise that their knowledge attributions are typically false because in everyday contexts knowledge attributions are made against a backdrop of shared assumptions about which possibilities are taken for granted in the present context. Bonjour (2010) argues that only sceptical infallibilism resolves the threshold problem, and later goes on to explain error in terms of people's unfamiliarity with the subtle philosophical reasons for thinking that knowledge

attributions are typically false. Dodd (2011) argues that fallibilism (and by concession, anti-scepticism) is incompatible with plausible assumptions regarding epistemic modality, and then explains error in terms of the incompatibility typically going unnoticed. Applying error as damage control post-argument to salvage a theoretically plausible view is perhaps appropriate in traditional intension- or extension-first analyses of knowledge, on which the aim is primarily to produce a counter-example free analysis of the target concept given in terms of necessary and sufficient conditions. But it sits uneasily within a function-first analysis concerned with investigating knowledge from the perspective of its practical, rather than theoretical, virtues. The role of error here is different. The function-first analysis that I propose centres error by showing how these patterns of error enable knowledge to serve a practical function in spite of sceptical truth conditions. On this picture, error is central to the analysis of knowledge precisely because it is central to knowledge being able to serve its practical function. This is not to say that knowledge's function requires the truth conditions posited by the sceptic in order to function, nor that anti-scepticism is ruled out on pragmatic grounds. Rather, the point is that the New Apraxia objection fails to demonstrate that sceptical truth conditions are anathema to knowledge serving a function. Contrary to what many function-first epistemologists have argued, appeals to the function of knowledge do not cut either for or against scepticism.

2.3.1. Strength of Epistemic Position

The first error that I discuss is widely posited by sceptics and regards agents' beliefs about the strength of their own and others' epistemic positions (see for example Butchvarov 1970:57; Bonjour 2010:71, Unger 1975:64, Dodd 2011: 683). According to this view, knowledge concept users systematically overestimate the strength of their own, and others', epistemic position. They are aware that knowledge comes equipped with maximally demanding standards but mistakenly take themselves and others to meet these standards, and so mistakenly believe many knowledge attributions to be true. It is clear how such a concept could fulfil a practical function. When an agent meets the strength of epistemic position the flagging of which would be expedient – perhaps, the strength of epistemic position at which they are sufficiently reliable informants – then both the agent themselves and those in their epistemic community mistakenly believe that the agent's epistemic position meets maximally demanding standards, and attribute knowledge accordingly. Insofar as flagging this epistemic position is what makes knowledge functional, then the sceptic's knowledge concept is also functional.

Peter Unger's (1975) famous defence of scepticism based on 'knowledge' being an absolute term can be viewed as an example of this strategy. For Unger, knowledge is useful because sentences of the form "S knows that P", although false, have true entailments, such as 'S is in a sufficiently strong epistemic position with regards to P for present purposes'. It is these true entailments that we have a practical need to communicate and further, for our practical purposes, it makes no difference whether "S knows that P" is true or only the entailments are true. Here is how Unger (1975:52, my italicisation) puts the point:

[E]ven if you do not really know, still, it may be that, for practical purposes, you are in a position with respect to the matter [...] which is not importantly different from

knowing. If this is so, then it may be better, *practically speaking*, for you to falsely believe that you know than to have no belief at all

In emphasising that having the concept of knowledge is practically better than having no concept at all, Unger offers exactly what the function-first philosopher is after: an explanation for the emergence and continued use of a concept that draws on the concept's ability to perform a given function, namely, to enable the communication of entailed propositions. Crucially, our practical needs to communicate the true entailments do not push the standards for knowledge down and away from those posited by the sceptic, since the sceptic's knowledge concept enables the communication of the true entailments just as well as an anti-sceptic's knowledge concept would. Although Unger does not specify what the second-order function of communicating the epistemic position that we find ourselves in is, his argument is amenable to being combined with any function-first hypothesis. One could argue, for instance, that communicating facts about one's epistemic position is practically important because it enables us to flag reliable informants.

Note that the function-first anti-sceptic cannot respond that even if the sceptic's knowledge concept is more useful than no concept at all, anti-scepticism is still to be preferred because the anti-sceptic's knowledge concept is more useful than both the sceptic's knowledge concept and having no knowledge concept at all. The function-first philosopher is only entitled to assume that our concepts are *adequate* for serving their function, not that they are *optimal* for serving that function. So perhaps the anti-sceptic's knowledge concept would serve knowledge's function better than the sceptic's knowledge concept, but this does not give us a reason to think that we have ended up with the better, anti-sceptic concept. This is the spirit that Unger's argument should be read in: the sceptic's knowledge concept would, from a practical perspective, work just as well as an anti-sceptical one. Centring the practical roles of knowledge does not, then, rule out scepticism.

We can build an even more explicit answer to the double bookkeeping objection inspired by Unger's remarks. Unger's entailment story aims to make plausible the claim that knowledge concept users typically overestimate the strength of their epistemic position: we fail to notice the difference between our actual epistemic position and the position demanded by the standards for knowledge because the difference is irrelevant from a practical perspective. A way of explaining why people overestimate the strength of their epistemic position is that they treat farfetched possibilities of error as if they were already ruled out. Function-first epistemology aims to analyse knowledge as functional in light of human needs, capacities and perspectives. Knowledge has been shaped by contingent facts about 'creatures like us', where these include "facts about our physical environment, our social organization, our cognitive capacities, and the basic aims and interests humans typically have" (Hannon 2019b:14); about "human psychology and the structure of human societies" (Kusch and McKenna 2020:1073); and about "the needs and capacities of concept-users and their environment" (Queloz 2021:29). If a contingent fact about 'creatures like us' is that we do not take sceptical possibilities seriously - if "our mind in a sense rules them out at once" (Craig 1990:10) - then we have no practical need for a knowledge concept on which knowing is compatible with such possibilities. Precisely because we automatically treat sceptical possibilities as having already been ruled out, we do not need our

knowledge concept to explicitly allow a knowledge attribution to be true even if there are uneliminated sceptical error possibilities.

I here understand ‘taking an error possibility seriously’ as not merely to entertain it, but to treat it as an error possibility which is “credible” (Pritchard 2015: 439); which is a “mistake you can actually make in this situation” (Stone 2000: 530); which is a “genuine possibility” (Butchvarov 1970:57); or which is a “live option” (Blome-Tillmann 2009:255). Taking a possibility seriously is to believe that the possibility could well obtain. There are several reasons why agents might not take sceptical possibilities seriously. The most straightforward is that they are not entertained in the first place, perhaps on account of their “complicated, subtle, and unobvious” (Bonjour 2010:72) nature. With what DeRose calls the “Aw, come on!” (2004:1) response to sceptical scenarios, even when entertained some possibilities are regarded as too improbable, farfetched or remote to engender serious consideration.

Indeed, that people in fact tend not to take remote error possibilities seriously is a claim that sceptics and anti-sceptics alike agree on. This need not, as the function-first anti-sceptic argues, shape the truth conditions of knowledge towards anti-scepticism, on which “S knows that P” is true even if there are uneliminated possibilities in which $\sim P$. Rather, it just as plausibly shapes the truth conditions for knowledge towards maximally demanding ones. However, when attributing knowledge in everyday life for the purposes of recommending reliable informants, possibilities which are not taken seriously are treated as if they have been ruled out. Double bookkeeping ensues: the attribution conditions for sentences of the form “S knows that P” are responsive to whether S has ruled out every $\sim P$; possibility that S (or the attributor, or the community) takes seriously. Meanwhile, the truth conditions for sentences of the form “S knows that P” are responsive to whether S has ruled out every $\sim P$ possibility. This version of scepticism diagnoses the frequent overestimating of our own and other’s epistemic position as due to the possibilities which we automatically fail to take seriously being treated as having already been ruled out. As in Unger’s original argument, the discrepancy between knowing and our actual epistemic position only rarely comes to our attention precisely because in most circumstances the difference between the two does not impede knowledge from fulfilling its practical purposes. Therefore, on the plausible assumption shared by sceptics and anti-sceptics that sceptical hypotheses are not generally taken seriously, the sceptic’s knowledge concept is equally as well suited to fulfilling knowledge’s functions.

2.3.2. Communicated Content

Another error that is frequently posited by sceptics and which explains double bookkeeping is error regarding the content typically communicated by knowledge attributions (see Stroud 1984, Stone 2000, Schaffer 2004, Chung 2018, Russell 2021, Stoutenburg 2021, Kyriacou 2020). These sceptics hold that knowledge attributions of the form “S knows that P” are not typically used to communicate that “S knows that P”. Rather, they communicate some other content, such as “S is in a good enough epistemic position with regards to P as to be regarded as a reliable informant as to whether P”. Taken at face-value, such knowledge attributions are typically false: they communicate that “S has ruled out every error possibility in which $\sim P$ ”, which is false in most of

the contexts in which it is uttered. However, these sceptics argue, the content that such sentences typically function to communicate is different, and often true.

To see how this works, consider some uncontroversial examples of sentences which typically communicate a content other than their face-value content: “I’m freezing” to communicate that one is (merely) very cold, “Marie arrived at 9” to communicate that Marie arrived approximately at 9 and “Jim is a bulldozer” to communicate that Jim is extremely pushy. These sentences are examples of exaggeration, loose talk, and metaphor respectively. What all of these sentences have in common is a bifurcation between the content communicated by a face-value (or ‘literal’) reading, and the non-face-value (or ‘non-literal’) content which the sentences are usually employed to communicate. Say that the content communicated by the face-value reading of these sentences (that I’m freezing, that Marie arrived at 9, that Jim is a bulldozer,) is false – still, the content that they are being employed to communicate (that I am very cold, that Marie arrived approximately at 9, that Jim is extremely pushy) may well be true.

That some sentences typically communicate a content other than their face-value content, then, is uncontroversial. Some sceptics hold that knowledge attributions belong in this category: sentences of the form “S knows that P” are not usually used to communicate their typically false face-value content, but some other, typically true, communicated content. Unlike the examples of exaggeration, loose talk, and metaphor given above, the sceptic posits the following error: we do not realise that knowledge attributions have a non-face-value reading which is typically communicated, and which differs from its face value reading. Whereas the ‘strength of epistemic position error’ sceptic posited double bookkeeping between the truth and attribution conditions for sentences of the form “S knows that P”, this strand of scepticism posits double bookkeeping between the content communicated by a face-value and non-face-value reading of sentences of the form “S knows that P”.

Jim Stone’s (2000) scepticism is an example of scepticism that posits error regarding what knowledge attributions communicate. On Stone’s view, a face-value reading of “S knows that P” is that ‘S has ruled out every $\sim P$ world’. Typically, the content communicated by these face-value readings is false. The communicated content, however, contains an implicit conditional operator and communicates something along the lines of: “supposing not-P worlds that conflict with my theory of the world are already ruled out, then, given E, it is indubitably certain that P” (2000:529) where E is one’s total body of evidence and one’s ‘theory of the world’ consists in assumptions the falsehood of which one does not seriously entertain as possible. Stone’s view double bookkeeps in the following way: the truth and attribution conditions for face-value readings of “S knows that P” come apart from the truth and attribution conditions for the communicate content of “S knows that P” because the latter includes an implicit conditional operator. Stone’s view is a form of scepticism that posits widespread error, namely, that people do not realise that their knowledge attributions typically communicate a conditional claim. Stone’s explanation for the error is reminiscent of Unger’s: given the shared assumption that farfetched error possibilities are not to be taken seriously and therefore may be treated as having already been ruled out, for all practical purposes there is no difference between conditional and non-conditional knowledge attributions. Whereas Unger takes this to support the existence of a factual error regarding the

strength of our epistemic position, Stone takes it to support the existence of an error regarding what our knowledge attributions typically communicate.

Error regarding the content communicated by knowledge attributions yields double bookkeeping, and also yields a sceptic's concept of knowledge that is capable of fulfilling the functions of knowledge. On this view, both the truth conditions and attribution conditions for knowledge bifurcate depending on whether "S knows that P" is given a face-value or non-face-value reading. On a face-value reading, the truth and attribution conditions are those posited by the sceptic, while on the more typical non-face-value reading, they are those posited by the anti-sceptic and include an implicit conditional operator. Knowledge is functional in virtue of its non-face-value reading. Knowledge is able to fulfil the function of recommending informants, for example, in virtue of the reading on which the sentence contains an implicit conditional operator ruling out sceptical scenarios. That the non-face-value reading is instrumental in knowledge serving its functions does not speak against scepticism – it means, rather, that the sceptic's knowledge concept describes the less ubiquitous of the two readings for sentences in which 'knowledge' appears. Crucially, however, this type of sceptic posits that everyday speakers do not realise that 'knowledge' has a non-face-value reading, and mistakenly take themselves to be making face-value knowledge attributions. Speakers who utter sentences of the form "S knows that P" mistakenly believe themselves to making a face-value claim that S knows that P, and mistakenly believe that this face-value claim contains no conditional operators. In fact, and unbeknownst to them, the non-face-value reading that their utterance communicate is something akin to "supposing that the not-P worlds that conflict with S's theory of the world are already ruled out, then, given E, it is indubitably certain for S that P".

The function-first anti-sceptic might nevertheless object that on this view the non-face-value reading is what enables knowledge to fulfil its practical roles, still leaving it mysterious why knowledge comes equipped not only with a different face-value reading, but specifically with the face-value reading posited by the sceptic. In response, recall first that the function-first sceptic need not show that her knowledge concept is more functional than the anti-sceptic's, but merely that it is functional. So long as knowledge admitting of different face-value and non-face-value readings enables speakers to recommend informants, answering the New Apraxia objection does not require the sceptic to additionally argue that coming equipped with the face-value reading posited by the sceptic renders knowledge more functional than if it did not come so equipped.

Second, the sceptical infallibilist may strive to make plausible the claim that positing bifurcation between face-value and non-face-value readings coheres well with the functionality of knowledge by pointing to the well-researched functions of non-face-value talk in general. An implicit conditional operator is only one of several candidate mechanisms that might explain how knowledge attributions come to communicate a different content, with different truth conditions, to that communicated on a face-value reading. Other mechanisms include exaggeration (Carter 1995), hyperbole (Schaffer 2004), metaphor (Chung, 2018, 2020), loose talk (Russell 2021, Dinges 2021) and Gricean implicature (Stoutenburg 2021). Each of these mechanisms is employed in everyday speech because doing so confers advantages or enables speakers to communicate something that they could not (or not so easily) communicate using exclusively face-value language. For example, metaphor functions to assist reasoning and inference making

by presenting novel and complex information in more familiar and simpler guises, thereby reducing cognitive load (Landau et al 2010, Johnson 2010, Bougher 2012, Thibodeau et al 2017), while exaggeration and hyperbole enable speakers to expressively convey their evaluation of a state of affairs (Ruiz de Mendoza and Galera 2014, Cartson and Wearing 2011, 2015, Popa-Wyatt 2020)²⁶. The function-first sceptic, then, can look to the functionality of her posited type of non-face-value talk for avenues by which to make plausible the claim that coming equipped with a non-face-value reading aids the functionality of knowledge²⁷.

2.4. Conclusion

The New Apraxia objection holds that the truth of sceptical infallibilism is incompatible with the plausible assumption that knowledge serves a practical function. In response, I have argued that the sceptical infallibilist can, and typically does, explain how knowledge serves a function in spite of its sceptical truth conditions by appeal to systematic patterns of error, and I have highlighted two different types of error that the sceptical infallibilist might posit to explain how false knowledge attributions might nevertheless be functional. The remainder of this thesis outlines and defends a particular way of divorcing communicated content from face-value content: fictionalism about knowledge, on which the sentences of the form “S knows that P” admit of different face-value and non-face-value readings.

²⁶ The proponent of the New Apraxia objection might find such widespread error implausible, especially concerning the meaning of sentences, a question on which we might expect speakers to have a high degree of introspective access (Liggins 2010). Chapter 5 addresses this objection in detail. For now, it’s worth noting that epistemologists have argued that other widely endorsed theories on knowledge such as contextualism (Hawthorne 2004:107, Stanley 2005:116) and non-sceptical invariantism (Montminy 2009:653, Kyriacou 2020:557) also posit semantic blindness of some form.

²⁷ Chapter 4 will defend the claim that there are advantages to the sceptical infallibilist’s concept of knowledge. See also Aikin (2019) and Kyriacou (2020) for practical advantages of knowledge concepts that entail scepticism.

Chapter 3: Fiction and Function

The previous chapter argued that compatibly with the concept of knowledge serving many functions, the semantic content of sentences of the form “S knows that P” might be divorced from the meaning that they are standardly used to communicate. This chapter examines in detail an example of such a view: Julianne Chung’s (2018, 2020) version of epistemic fictionalism, according to which knowledge talk is typically non-literal talk. The thesis will go on to develop and defend a related way of being an epistemic fictionalist. As the only example of a fictionalist view about knowledge in the literature, examining Chung’s view will serve as a springboard from which to position the version of fictionalism that I prefer.

First, I introduce fictionalism in general, and discuss five criteria which a fictionalist theory should meet if it is to be well-motivated (§3.1.). I highlight some arguments typically taken to support fictionalism (§3.2.), and then outline Julianne Chung’s version of epistemic fictionalism (§3.3). I argue that despite its many virtues, it fails to meet one of the five criteria for a well-motivated fictionalist theory (§3.4.) Finally, I consider a refinement of Chung’s original fictionalist theory, and argue that although it constitutes an improvement, it still doesn’t resolve the problem faced by the original account (§3.5.).

3.1. Fictionalism in Philosophy

3.1.1. Fictionalism: a first characterisation

In their introduction to *Fictionalism in Philosophy*, Armour-Garb and Kroon (2020:1) write that “[i]t is easier to feel the force of fictionalism in philosophy if one begins by considering a problem that emerges from some of what we say and appear to believe”. This is no doubt why most discussions of fictionalism begin with a puzzle about some domain of discourse, and this thesis is no exception. Here, then, is a puzzle about mathematical talk: talk of there being 12 apostles, 250 grams of sugar, and infinitely many primes appears to commit speakers to the existence of abstract platonic entities – in this case, to the existence of the number 12, the number 250 and prime numbers. Yet, many speakers who utter such number sentences do not, or would not wish upon reflection, to commit themselves to the existence of abstract platonic entities called ‘numbers’. Put differently, there is a mismatch between what such sentences, when taken at face-value, appear to commit us to and the way that we take the world to be – a mismatch which suggests that the sentences are, perhaps, untrue. Or at least, that speakers who utter them would, upon reflection, typically consider them to be untrue. Nevertheless, the suggestion that since we do not believe in abstract platonic entities we ought to cease using number talk in all domains, from cooking to physics, would be met with incredulity and impatience. Number talk is simply too useful in both everyday and theoretical contexts to do without. Furthermore, when using number talk in the course of daily life the question of whether abstract platonic entities exist simply does not strike us as important or relevant. This presents a problem: on the one hand, mathematical talk is pervasive and indispensable; on the other hand, it appears to commit users of the discourse to widespread error, and does not reflect speakers’ deeper commitments regarding what the world is like.

Stephen Yablo (1998, 2000) proposes a disarmingly simple explanation for this apparent tension: mathematical talk does not commit us to the existence of abstract platonic entities because it does not typically communicate the content given on a face-value reading. In other words, mathematical talk is non-literal. Just as the sentence “Juliet is the sun” does not commit Romeo to the claim that Juliet is a blazing ball of plasma around which the Earth rotates, so the sentence “There are 12 apostles” does not commit speakers to the claim that there are abstract platonic entities. In general, uttering non-literal sentences does not commit speakers to the truth of the target sentences’ face-value semantic content. Rather, competent interlocutors understand speakers to be aiming to achieve something other than communicating the face-value meaning of their non-face-value sentences. A mother who utters the sentence “Once upon a time there was a dragon who lived in a castle” in the course of telling a story, for example, would be taken to be aiming to entertain. A child who utters the sentence “We’re the Clangers” would be taken to be aiming to initiate a game of make-believe. And Romeo uttering the metaphorical sentence “Juliet is the sun” should be taken as trying to communicate some other content that he takes to be true, but for which he perhaps lacks the face-value language.

The fictionalist says that, analogously, speakers who utter sentences in their target domain of discourse do not (or perhaps, should not) be interpreted as aiming to communicate the face-value meaning of their sentence but rather to do something else entirely. To be a fictionalist about mathematical discourse is to say that number sentences such as “There were 12 apostles” do not communicate their face-value content (roughly, that there are abstract platonic entities called numbers and the amount of apostles corresponds to the abstract platonic number ‘twelve’) but communicate some non-face-value content (perhaps, that *if* there are abstract platonic entities called numbers, *then* the amount of apostles would correspond to the abstract platonic number ‘twelve’). It is this latter content that agents typically wish to communicate when they utter mathematical sentences. Whether sentences in mathematical talk are true or false on a face-value reading, then, does not determine whether the content that speakers typically aim to communicate when they utter such sentences is true or false. In this way, the mathematical fictionalist tries to dissolve the apparent tension between what number-talk appears to commit speakers to and what speakers take the world to be like: number-talk does not aim to communicate its face-value reading, and therefore, it does not commit speakers to anything that they might repudiate.

This is the core spirit of fictionalism. At their heart, fictionalist theories posit that sentences in their target discourse have a ‘face-value content’: the surface grammatical and semantic meaning of a sentence (e.g. ‘Juliet is a blazing ball of plasma’, ‘There are abstract platonic entities’). Yet, the main aim of uttering these sentences is not to communicate this face-value content. Most fictionalist theories posit that there is a ‘non-face-value’ content, which designates the content that the non-literal sentence actually communicates (e.g. ‘Juliet is radiant and worthy of worship’, ‘if there are abstract platonic entities called numbers, then the amount of apostles would correspond to the abstract platonic number ‘twelve’). Roughly, ‘face-value’ content tracks the content sometimes referred to as ‘literal’ while ‘non-face-value’ content tracks the content sometimes

referred to as ‘non-literal’, ‘fictional’ or ‘expressed’²⁸. Thus fictionalist theories are linguistic theses: they are theses about fragments of discourse, rather than the real world phenomena that these fragments of discourse purport to refer to (if indeed, they refer to phenomena at all). As Armour-Garb and Woodbridge (2015:5) stress, “the fact that fiction is a category of discourse means that one can apply the notion of fiction only to discourse”. Strictly speaking, the moral fictionalist does not claim that moral facts are fictional; she claims that sentences of the form “it is immoral to ϕ ” are fictional. In the same way, a fictionalist about knowledge does not claim that knowledge is fictional, but rather that sentences of the form “S knows that P” are fictional.

Fictionalism has received increased attention in recent years, and philosophers have argued that fictionalism is the best account of a wide range of discourse domains, including morality (Joyce 2005, Kalderon 2005), mathematics (Yablo 2010, Marcus 2015, Balaguer 2009, Bueno 2009), some aspects of scientific talk (Rouse, 2008, Levy 2011), modelling (Godfrey-Smith 2007, Frigg 2008), objects (Van Inwagen, 1977), present objects (Miller 2021), composite objects (Rosen and Dorr 2002), existence (Armour-Garb and Woodbridge 2015), truth (Burgess and Burgess, 2011, Armour-Garb and Woodbridge 2015), grounding (Thomson 2022) fictional characters (Lewis 1978, Brock 2002), God (Eshleman 2005, Le Poidevin 2020, D’Amato 2013), law (Gawthorne, 2013), modality (Rosen, 1990) mental states (Sprevak 2013, Toon 2016), colours (Gatzia 2010), gender (Logue 2022) and knowledge (Chung 2018, 2020), among others.

Despite its increased popularity as an antirealist strategy, there is no standard way of understanding what fictionalism is. I now turn to three quick distinctions that fictionalists typically make, but which won’t concern the arguments in this thesis very much. First, in terminology introduced by Burgess (1983: 96), and popularised by Stanley (2001:36) in the context of fictionalism, it has become standard practice to distinguish between *hermeneutical* and *revolutionary* fictionalism. Hermeneutical fictionalism is a descriptive thesis, holding that as a matter of fact the target discourse under investigation is already being used to communicate non-face-value content. In contrast, revolutionary fictionalism is a prescriptive thesis, holding that we *ought* to start using a discourse domain to communicate non-face-value content. Say that sentences of the form “ ϕ is moral” actually are typically used to communicate their face-value meaning, and furthermore that such sentences are eminently useful; nevertheless, say that careful philosophical reflection has revealed that in spite of its ubiquity and usefulness, moral talk commits us to a problematic moral ontology. The revolutionary fictionalist might advocate that we start using moral talk in the way that the hermeneutic fictionalist suggests that we already do. That is to say, the revolutionary fictionalist recommends that we stop using the discourse in a face-value spirit, and start using it in a non-face-value way. This thesis concerns hermeneutic fictionalism exclusively, and hereafter the term ‘fictionalism’ refers to hermeneutic fictionalism. I shall not be concerned with the question of whether knowledge talk should be fictional, but on

²⁸ I use Armour-Garb and Woodbridge’s (2015:6) terms ‘face-value’ and ‘non-face-value’ rather than the terms ‘literal’ and ‘non-literal’ because it avoids the confusion between standard (as in, typical) and literal (as in, non-fictional) use. It also avoids the problem that some sentences arguably do not have a literal meaning.

whether it is already typically used fictionally to communicate something other than its face-value content.

The second distinction is between assertoric and non-assertoric fictionalist theories. Every fictionalist agrees that in uttering sentences in the target domain of discourse, speakers do not commit themselves to the truth of their uttered sentences' face-value meaning. But fictionalists disagree about what precisely one *does* commit themselves to instead, in uttering these sentences. Assertoric fictionalists think that in uttering a sentence in the target discourse domain, a speaker asserts (and believes) the non-face-value content of the sentence. Assertoric fictionalism is a semantic thesis, holding that semantic redirection occurs so that in uttering a non-face-value sentence, one comes to assert (and believe) its non-face-value content. Operator fictionalist theories such as Lewis' (1978) fictionalism about fictional characters is an example of this view, on which the non-face-value meaning of an utterance " ϕ " is "in the fiction f , ϕ ". In contrast, non-assertoric fictionalist theories think that in uttering a non-face-value sentence " ϕ ", the speaker does not assert anything at all. Rather, they engage in a different speech act altogether. This might be the speech act of quasi-asserting²⁹, pretending to assert³⁰, or launching a game of make-believe³¹. On this view, sentences of the form "it is immoral to ϕ " are more closely related to sentences of the form "Let's pretend that ϕ " - they might be paraphrased as "Let's pretend that it's immoral to ϕ " (Joyce 2022).

Something in the ballpark of this distinction is often cast with different terminology: what I have called assertoric fictionalism is similar to 'content' (Eklund 2019) or 'cognitive' (Joyce 2005:295) fictionalism, and what I have called non-assertoric fictionalism is similar to 'force' (Eklund 2009), 'non-cognitive' (Joyce 2005:295) or 'instrumental' (Yablo 2001:74) fictionalism. The content-force distinction is supposed to capture the divide between fictionalists who think that in uttering sentences in the target discourse domain a speaker succeeds in asserting some new non-face-value *content*, to contrast with fictionalist theories on which sentences in the domain of discourse are uttered with non-assertoric *force*. Although I will sometimes use the terminology of 'content' and 'force' fictionalism (especially in chapter 5), I prefer the assertoric and non-assertoric distinction because it more usefully categorises force fictionalist views on which a non-assertoric utterance of " ϕ " nevertheless succeeds in asserting some non-face-value content which the speaker believes. Toon's pretence theory about mental states is an example of such a view. On Toon's fictionalism, speakers who utter sentences of the form "R believes that P" are *pretending* to assert "R believes that P", and through the act of engaging in this pretence, nevertheless succeed in *asserting* that R is in some state S such that, fictionally, the speaker says something true³².

The third distinction concerns the role of attitudes. Some philosophers (e.g. Kalderon 2008, Daly 2008, Jenkins 2006) take as essential to fictionalism that speakers bear a certain attitude

²⁹ E.g. Rosen (1994), Daly (2008), Kroon (2008)

³⁰ E.g. Toon (2016), Liggins (2010), Miller (2021), Joyce (2022)

³¹ E.g. Walton (1993)

³² See also Walton (1993) on fictional characters and Armour-Garb and Woodbridge on existence and truth talk (2015).

towards the sentences in the target domain of discourse. On this view, to be a fictionalist is to say that users of the target domain of discourse do not believe the face-value content of their utterances. Rather, they merely accept³³, exploit³⁴, or entertain³⁵ the sentences in *D*, where these attitudes all fall short of belief. Thus Kalderon (2008: 40) writes that “[f]ictionalism is not primarily a semantic thesis. It is a claim about the nature of acceptance”. Daly echoes this sentiment, writing that “[f]ictionalism is an account of what attitude we should take to the sentences of [a subject matter] *S*. It says that we should not believe them” (2008:424). Other fictionalists opt for semantic theories which do not trade on the notion of acceptance at all. Armour-Garb and Woodbridge’s fictionalism about existence-talk, for example, posits a semantic redirection so that speakers do believe the content of their sentences – it is just that they believe the non-face-value content, as opposed to the face-value content. Furthermore, as Joyce (2005) points out, operator fictionalist views hold that sentences in their target discourse contain tacit operators, so there is no obstacle to speakers straightforwardly believing the content of their sentences. For example, sentences such as “there are infinite primes” contain a tacit operator: “[According to the fiction of mathematics] there are infinite primes”, which speakers might well believe (and which might well be true). Along with Joyce (2005) and Armour-Garb and Kroon (2020), I don’t wish to restrict fictionalism to a theory about attitude. Attitude-wise, what matters for a fictionalist view is that speakers do not believe the *face-value* content of their sentences. Compatibly with this, a view might count as fictionalist without positing any particular attitude towards sentences in the target domain of discourse.

Having said something about the distinctions that I don’t focus on, I now turn to two distinctions that do determine whether a fictionalist theory is well-motivated: fictionalism must explain the function of its target domain of discourse; and fictionalism must employ the notion of fiction to explain how sentences in its target discourse domain succeed in fulfilling their function. My main aim is to sketch a well-motivated fictionalism, a fictionalism that gives us a reason to think that the target discourse is indeed fictional³⁶.

3.1.2. *Functional Fictional Talk*

Essential to fictionalism is the claim that the domain of discourse given a fictionalist treatment is somehow useful. In this sense, fictionalism shares a key commitment with theories generated through a function-first analysis. We have seen how investigating a concept using a function-first lens means understanding every concept as fulfilling a purpose or a function. Many realists about a domain of discourse take the function of the concept to be tracking some natural fact about the world: the moral realist says that the purpose of moral talk is to track moral facts (Wielenberg 2010); the mathematical discourse realist will say that the purpose of moral talk is to track number facts (Clarke-Doane, 2012); and at least one realist about knowledge says that the purpose of knowledge talk is to track epistemic facts (Kornblith 2002). It is of course possible to be a realist

³³ Kalderon (2008)

³⁴ Daly (2008)

³⁵ Hills (1997)

³⁶ Although I make some comments about distinguishing fictionalism from other views, my main aim isn’t to taxonomise – I acknowledge that many views which are commonly described as fictionalist, both by their advocates and by others, do not adhere to these criteria.

who doesn't think that the function of a domain of discourse is just to track natural facts about the world. The function-first project is not typically conceived of as an antirealist project, but neither does it posit that the functions served by knowledge are tracking functions. Indeed, as Craig (1990) and Hannon (2019a) point out, even if epistemologists did succeed in the project of identifying a natural fact about the world which perfectly matches the extension of our concept of knowledge, they would not be out of a job: it would still be instructive to ask what it is about these particular standards that has been worth picking out with the concept of knowledge. As discussed, many function-first theorists build into their method that the resulting concept must be checked against intuitive extension and intension, where this means that the resulting concept should feature in sentences that are typically true – barring “some half-way plausible hypothesis about how the massive misfit could have come about” (Craig, 1990:71) that is.

Unlike typical theories generated on function-first analyses, fictionalist theories posit that in addition to being useful, a domain of discourse is somehow infelicitous – perhaps, infelicitous in the sense of being false. A key hallmark of fictionalism is to outline and defend the relevant “half-way plausible hypothesis”: to account for the function of a domain of discourse which is nevertheless infelicitous (and perhaps, which is nevertheless false). This is reflected in many characterisations of fictionalism. Vaihinger's (1911) Kantian inspired *Philosophy of as If* is often cited as the ancestor of modern fictionalism (see e.g. Fine 1993), and emphasises throughout the role that fictions play in our cognition. For Vaihinger, fiction selection is driven by considering which fiction is most expedient given the purposes of a domain of discourse. Most modern day fictionalists do not wish to follow Vaihinger in claiming that every useful concept is fictional, but do continue to define fictionalism as the thesis that a domain of discourse may be infelicitous at face-value while nevertheless being useful.

Here are some representative definitions of fictionalism that all reflect the idea that fictionalism treats its target discourse as useful, albeit (potentially) false:

Fictionalism about a discourse claims that the sentences of the discourse are useful but does not claim that they are true. Typically, fictionalists will deny that sentences of the discourse are true. (Liggins 2012)

Utterances of sentences of the discourse are best seen not as efforts to say what is literally true, but as useful fictions of some sort. (Eklund 2019)

Fictionalism proposes that a certain discourse – such as talk about possible worlds, or mathematical talk – is useful, perhaps even indispensable for theoretical purposes, but should not be taken as true. (Caddick Bourne 2018)

Fictionalism about region of discourse *D* claims that those participating in *D* should not have (and perhaps, in the case of certain types of discourse, typically do not have) truth as their aim when they accept a sentence from *D*. (Kroon 2011: 787)

The “fictionalist umbrella”, in that claims that are not *literally* true can still be usefully used for a variety of purposes. (Nolan 2005:2014)

Schematically, we can formulate fictionalism about S as follows: (1) the fictionalist does not believe any sentence of S, but (2) he believes that the sentences of S are useful. (Daly 2008:425)

Central to fictionalism, then, is the idea that whatever the utility of a discourse is, its utility must be independent of the truth of the sentences in the discourse. Toon (2016), for example, argues that when closely examined, talk about mental states such as belief and desire do not appear to reflect a genuine commitment that there are such states and that these states cause behaviour. Rather, we can understand mental state talk as usefully fictional talk, whose purpose is to help describe, explain and make inferences regarding people's behaviour. Joyce (2001) argues that moral talk serves the function of preventing us from succumbing to potentially costly weakness of will, and expands co-operation to members of the community that are not directly related. Yet, these function are served entirely independently of the truth of sentences of the form "S believes that P" or "it is morally permissible to ϕ ". In short, what makes fictional sentences useful is not their truth³⁷.

Notice that so formulated, the fictionalist does not claim that typical sentences in her target domain of discourse actually are false. Rather, it suffices that the sentences be somehow infelicitous (Armour-Garb and Woodbridge 2015)³⁸. One way in which sentences in D might be infelicitous is by being false in a manner that we would expect speakers to detect. Other ways of being infelicitous include being semantically, epistemically or pragmatically infelicitous: perhaps we have good reason to believe that sentences in D are typically false, or generate contradictions when used in different contexts, or their truth value is indeterminate from our perspective, or perhaps they are true but using them to utter their face-value truth would not serve the aims of the discourse.

There is a sense, then, in which fictionalist theories are natural allies with function-first accounts: both take as a starting point observations about a discourse's utility. The fictionalist, however, is motivated additionally by widespread infelicity surrounding a domain of discourse.

3.1.3. Putting the 'fiction' in 'Fictionalism'

Fictionalism defined in terms of a somehow infelicitous, but useful, discourse is still missing something important: a positive reason to think that sentences in the target domain of discourse are typically uttered in a fictional spirit. After all, many anti-realist theories, including error theories, pragmatist theories and expressivist theories, hold that at least on some readings a domain of discourse is infelicitous yet useful. Mackie (1977:239), for example, famously describes moral obligations as "useful fictions" – what he means by this is that sentences of the form "it is morally obligatory to ϕ " are useful, yet false. Moral discourse on this view might also

³⁷ Joyce makes these points in the course of defending revolutionary moral fictionalism, but they apply equally to hermeneutical fictionalism.

³⁸ Although there is no logical barrier to claiming that sentences in D are felicitous but fictional, this avenue is never pursued because without some infelicity in the discourse, fictionalism would be severely under-motivated. If a domain of discourse is functional and felicitous on a face-value reading, then we have no reason to suppose that such sentences do something other than communicate their face-value reading.

be infelicitous, at least insofar as ‘false according to philosophical theory about morality’ counts as a form of infelicity. But because Mackie thinks that speakers who utter moral sentences take themselves to be saying something true (perhaps they do not notice, or are not moved by, the infelicity), he is an error theorist as opposed to a fictionalist. Van Fraassen’s (1980) account of scientific discourse is frequently described as fictionalist³⁹. Van Fraassen holds that we should accept, rather than believe, our best scientific theories – but he makes no use of the notion of fiction to flesh out what acceptance means, or why we should accept, as oppose to believe, our scientific theories. Van Fraassen’s view, then, ought not qualify as fictionalist.

Bourne and Caddick Bourne (2020:168) write that “[w]hat characterizes a fictionalist approach to subject matter X is the suggestion that X can be understood by appeal to the notion of fiction. Otherwise, fictionalism does not deserve its name.” Yet for some fictionalists, all that a fictionalist requires to “deserve the name” is to posit some similarity between paradigmatic instances of non-face-value discourse, such as literary fiction, and the target domain of discourse. The similarity might be ontological, in the sense that both sentences in the target discourse domain and sentences in canonical examples of fiction make reference to entities that do not exist (or, to more accurately frame this point in terms of discourse, hold that both sentences in the target discourse domain and sentences in the canonical example of fiction are false if given a face-value reading). Thus Sainsbury (2009:175) writes “a theory is fictionalist concerning some region of thought if it claims that regions are, or can be, or should be regarded as possessing this similarity with fiction: they don’t have to be true to be good”. The similarity might be epistemic, in the sense that both sentences in the target discourse domain and sentences in the canonical example of fiction are not typically believed, but are merely accepted or exploited. In the SEP entry on modal fictionalism, Nolan (2016) invokes this analogy in his treatment of fictionalism, characterising it as “an approach to theoretical matters in a given area which treats the claims in that area as being in some sense analogous to fictional claims: claims we do not literally accept at face value, but which we nevertheless think serve some useful function.” The similarity might involve the function of the discourse. For example, Joyce (2013:520) describes fictionalism as the “thesis that the role that these [non-face-value] utterances do play is substantively similar to the role of familiar fictional discourse”.

There are several problems with defining fictionalism purely as the thesis that one’s domain of discourse must somehow be similar to a canonical example of fiction. The first is simply a matter of classification. As Burgess (2004:20, original italics) puts it, “it is impossible to quarrel with the proposition that mathematics is *in some respects* like fiction. For indeed, *anything* is like *anything* else, *in some respect*”. Mere similarity is too coarse-grained a tool to usefully distinguish fictionalism from other views – indeed, even some realist views might hold that fragments of domains of discourse are in some respects similar to fiction. For example, any realist about fictional entities will hold this position just in virtue of their subject matter, and so would trivially count as fictionalist. And other anti-realist theories such as pragmatism and expressivism also posit that the target domains of discourse share a similarity with paradigmatic instances of literary fiction, for example, that sentences in the target discourse domain are accepted or quasi-asserted

³⁹ See for example Field (1980), Yablo (2001), Hussein (2004), Kalderon (2005), Rosen (2005), Blackburn (2005), Nolan et al. (2005), Daly (2008), Sainsbury (2010) and Eklund (2011).

rather than believed and asserted. For example, Blackburn's (1984, 1988) quasi-realism posits that moral sentences are quasi-asserted, in the sense that they should be heard not as asserting their content, but in expressing certain attitudes. Yet as Blackburn (1987) makes clear, quasi-realism is not intended to be a type of fictionalism: it makes no reference to notions of fiction, and holds that sentences in its target discourse domains are truth-normed⁴⁰.

Another limitation of characterising fictionalism in terms of similarity to paradigmatic instances of fiction is that it demands some set of distinctive features constitutive of fiction to which features of the target discourse are relevantly similar to. Take Kalderon's claim that fictionalism is characterised by similarity in attitude between moral talk and canonical examples of fictional talk, such as literary fiction. Not only does this saddle the fictionalist with what Armour-Garb and Woodbridge (2015:11) term the 'added burden problem' of needing to commit to a particular philosophical view on which attitudes are typical of engagement with fiction, but it assumes that there is one attitude involved in familiar fictional discourse. The problem with this line of reasoning is that there might not be a well-defined demarcation between fiction and other types of representation or set of characteristics, such as attitudes, shared by all instances of fiction (Friend 2012, Bourne and Caddick Bourne 2020).

A more substantive problem with characterising fictionalism in terms of similarity between the target discourse and paradigmatic examples of fiction, however, has to do with motivation: pointing to similarity gives us no reason to think that the target discourse itself operates along the lines of the suggested similar fiction. Say that, following Joyce (2022), one's target discourse plays a similar role to the role played by a given fictionalist discourse. Where does the notion of fiction come in? Joyce would first need to show that the domain of discourse indeed fulfils this particular role; independently argue that a given type of fiction also fulfils this role; and finally, argue that these roles are relevantly similar. But what is the point of the comparison with paradigmatic examples of fiction? If the domain of discourse actually does fulfil that role, it would do so even if no instance of fiction fulfilled that role. Indeed, the target discourse may well have fulfilled this role even if there were no such thing as fictional discourse. Whether or not some types of fiction also play the posited role has absolutely no bearing on whether *this* domain of discourse fulfils the posited role. So on this picture, the notion of fiction is inert in the explanation of how one's discourse operates. Fiction enters the scene after the analysis of the target discourse has already been completed. Noting similarities in role would not, for example, suggest that there are any further similarities between the two types of discourse, nor would it give reason to think that the target discourse as a whole is fictional.

This point is not merely that this way of classifying fictionalism makes analysing whether a fragment discourse is fictional redundant, but more importantly that it undercuts the motivation for being a fictionalist about a domain of discourse in the first place. On this picture, one can count as a fictionalist even if they take there to be absolutely no reason to think that the fragment of discourse involves fiction of any type: all they require is that the discourse is in some sense similar to fictional discourse. In contrast, a philosopher who argued that similarities give us reason

⁴⁰ Although, see Lewis (2005) for the argument that quasi-realism is a type of fictionalism; see Blackburn (2005) and Jenkins (2006) for replies.

to think that the target discourse is itself an example of that type of fictional discourse would be well-motivated. For example, someone who thinks that a domain of discourse fulfilling the same role as metaphor makes it plausible that the target discourse *is* metaphorical is not subject to this criticism, nor is someone who thinks that a domain of discourse involving the same attitude as telling a story telling gives us reason to think that the domain of discourse *is* an instance of story-telling. Such a philosopher would be using the notions of metaphor and story-telling as integral to their analysis of the target discourse⁴¹.

This suggests that a more motivated version of fictionalism employs the notion of fiction as part the *explanation* for how the discourse succeeds in fulfilling its useful roles. That is to say, to really motivate the claim that a domain of discourse should be analysed along the lines of a fiction, the fictionalist needs to explain how the notion of fiction enables sentences in the target domain to fulfil the function in virtue of which they are useful (and hence, in virtue of which they have remained in circulation despite their infelicity). Looking back to the function-first method introduced in Chapter 2 brings into focus a more motivated way of incorporating fiction back into fictionalism. Recall that the key insight of the function-first method is to begin with an account of what a concept or fragment of discourse does, and work backwards to uncover its intension and extension by reflecting on what a concept would need to look like in order to fulfil these functions. I also argued that where the posited function is not merely to track a natural fact about the world, we should not assume that a function-first analysis will yield a realist picture of the concept being analysed. I think that function-first can do even more work for the fictionalist: running a function-first analysis can potentially reveal where a domain of discourse is likely to be fictional, by revealing how a domain of discourse could best serve its purpose if it were fictional. And the way in which a function-first account will reveal this is by showing how discourse involving the concept could not (or could not as well, or as easily) serve its function if it were not fictional.

To illustrate, take Yablo's canonical mathematical fictionalism. Yablo writes that once we grant that there is such a thing as metaphorical content, "then here is what we want explained: what are the reasons for accessing them *metaphorically*?" (1998:250). Yablo goes on to argue that certain types of content can only be accessed metaphorically. This is because there simply might not be a face-value way of communicating them; they might not be easily comprehensible when presented at face-value; or perhaps the face-value way of communicating them is under-determined. In the case of number-talk, speaking metaphorically enables speakers to utter sentences of the form "the As and Bs are equinumerous" even though it is indeterminate from their perspective whether abstract platonic entities exist – and hence the face-value reading of the utterance is true, and one that speakers would be happy to be taken to have said – or whether they do not, in which case speakers would want to be taken to have uttered a non-face-value content. Yablo's insight is to bake the indeterminacy into the discourse itself. He does this through pretence-style metaphor: metaphorical talk enables the communication of indeterminate content. Indeed, Yablo is explicit in taking his fictionalism to fit at least part of what a Carnapian explication is supposed to generate: a "well-disciplined, practically advantageous way of talking

⁴¹ See e.g. Yablo (2010) and Chung (2018) for accounts that argue from similarities to fiction to the conclusion that the target discourse domain is fictional.

that makes no pretence of being ‘really true’” (1998:242). A function-first account of number talk thus has the potential to lend support to this fictionalist analysis, by embedding the sense in which the discourse needs to be fictional to be functional into a more general quest to investigate number-talk by thinking about its functionality.

I argued in chapter 2 that the function-first philosopher always runs the risk that useful error is confounding her analysis. The same is true of fictionalist interpretations. It could be that particular patterns of useful error make it such that speakers using number-talk really do believe the face-value claims that they make, and yet their number-talk fulfils its useful function; it could be that users of number talk are simply mistaken about what their ontological number commitments are. We can never rule out error. Nevertheless, if it turns out that certain useful content appears to serve a function, or serve it much better, if presented in a non-face-value guise, then this gives the function-first philosopher at least a *prima facie* reason for thinking that the target discourse is best understood as fictional. In addition to intension and extension, we can apply the function-first method to reveal a concept’s ‘*face-value-ness*’: whether it is typically used in face-value or non-face-value sentences, by appeal to whether a non-face-value use would increase functionality.

A fictionalist account, then, applies the notion of fiction to the target discourse domain in order to explain how it is that the discourse succeeds in fulfils its central function, despite being infelicitous at face-value. This might take the form of showing how pretence is integral to how a discourse succeeds in communicating its non-face-value content, where the utility of the discourse lies in communicating this non-face-value content. It might take the form of applying some aspects of the semantics, epistemology or pragmatics of a particular type of fictional discourse to the target discourse domain, on the basis that these are the semantics, epistemology or pragmatics that would enable it to fulfil its function. Or it might involve arguing that the function of the domain of discourse is the same as the function fulfilled by a particular type of fictional discourse, giving us reason to think that the discourse domain belongs to that same type¹². Crucially, keeping an eye trained on the function of the discourse and working backwards to reveal its contours, as the function-first philosopher recommends, is a useful tool for revealing whether those contours contain non-face-value content.

3.1.4. Five Criteria

To summarise, I have argued that a well-motivated fictionalist theory:

- i. is a theory about sentences in a domain of discourse **D**
- ii. holds that sentences in **D** would be infelicitous on a face-value reading
- iii. holds that uttering sentences in **D** fulfils a useful function
- iv. holds that uttering sentences in **D** fulfils a function in virtue of doing something other than assert their face-value content

¹² See Armour-Garb and Woodbridge (2015) and Armour-Garb and Kroon (2020) for similar suggestions, not framed in terms of function-first philosophy. Both sets of authors limit fictionalism to views that involve some notion of fiction in the explanation for how sentences in the target discourse domain come to communicate their non-face-value content.

- v. applies the notion of fiction in the explanation of (iv), i.e. the notion of fiction explains how uttering sentences in D does something other than assert their face-value content

The aim is not (mainly) to taxonomise: I do not wish to insist that any theory that does not adhere to these criteria cannot properly call itself fictionalism. Rather, I hope to have defended the idea that any fictionalist theory that does not adhere to these is *under-motivated*. That is to say, unless a fictionalist can show that her target discourse meets these criteria, we do not have sufficient reason to think that the discourse in question is fictional after all. The remainder of this thesis will employ this characterisation of fictionalism.

3.2. Some Arguments for Fictionalism

Having said something about what fictionalism is, I briefly outline three key motivations for fictionalism, including an argument that I already started discussing in §3.1.3., and which will be particularly relevant to epistemic fictionalism: arguments from the function of fiction.

3.2.1. Arguments from Ontology

Historically, ontological arguments for fictionalism have been favoured, according to which we have good philosophical arguments to suggest that a given class of entity does not exist, and so good reason to re-analyse talk that is ostensibly about that entity so that it comes out not committing everyday speakers to its existence. Ontological arguments are, according to Yablo (2010:193), “dead and gone” – and for good reason. To work, ontological arguments need to claim that everyday talk is shaped by ontology: by what philosophers have good philosophical reason to think does or does not exist. But this strikes many philosophers as “comically immodest” (Burgess 2002:30). Lewis (1991:59), for example, writes that he’s “moved to laughter at the thought of how *presumptuous* it would be to reject mathematics for philosophical reasons”⁴³. Burgess and Lewis find fictionalism comedic because they take mathematics to be a more reliable source on the question of whether there are numbers than philosophy is. I agree with Daly (2006) that this argument is somewhat misguided – mathematicians may be experts on mathematical questions concerning numbers, without being experts on philosophical questions concerning numbers. But the more pertinent question is: why should we expect ontology to shape our concepts and discourse, when discovering the truth regarding this ontology demands complex philosophical argumentation? Barring some account on which a concept serves the specific function of tracking a natural fact, we are not entitled to suppose that our concepts to reflect reality in this way. Rather, as I argued in Chapter 2, we should expect our concepts to have the shape that they need to serve their function, given facts about us and the use to which we put them. Since few concepts have been shaped against a backdrop of philosophical use, we have no reason to think that our concepts reflect a complex ontology. For this reason, ontological arguments can be put to the side.

⁴³ See also Colyvan (2001) and Shapiro (1997) for similar misgivings; see Daly (2006) for a response.

3.2.1. Arguments from Function

I argued above that well motivated fictionalist theories centre the notion of fiction as part of the explanation for why we have the discourse to begin with, and how it serves its purposes despite being false at face-value. A good way of arguing for the claim that a domain of discourse is fictional stems from arguing that it being fictional (e.g., it communicating a non-face-value content) enables us to achieve something that we could not, or could not so easily achieve, without the discourse being non-face-value.

Yablo identifies three functions played by paradigmatic cases of non-face-value discourse, and suggests that either one, or perhaps a mixture of all three, can explain the role of non-face-value number talk. Non-face-value discourse might be:

1. Representationally essential, enabling us to represent content for which there is no direct face-value paraphrase (e.g. there is no face-value paraphrase for ‘angry clouds’ and ‘piercing glances’)
2. Presentationally essential, enabling us to reason about propositions in a way that affords greater cognitive advantages (e.g. ‘the average planet has 2.4 stars’ is much easier to process than its face-value equivalent)
3. Procedurally essential, enabling speakers to communicate a truth even if they are uncertain about whether the truth they wish to communicate is true on a face-value or non-face-value reading (e.g. ‘Jane was blessed with a speedy recovery’, where the speaker is agnostic about the spiritual ontology necessary for a face-value blessing)

Yablo goes on to argue that mathematical discourse is best understood as either presentationally or procedurally essential (or perhaps a combination of the two). Sentences of the form “there are 12 apostles” present a proposition that would otherwise be difficult to grasp (perhaps: ‘there is an apostle, and an apostle, and an apostle, etc.’) in a more cognitively accessible manner. Alternatively, perhaps whether there are abstract platonic numbers is indeterminate from speakers’ perspectives. In these cases the expression is intended to convey *either* that abstract platonic numbers exist and the platonic number twelve describes the number of apostles, *or* the paraphrase ‘there is an apostle, and an apostle, and an apostle, etc.’ Additionally, Yablo argues that the practice of using non-face-value discourse to achieve practical goals is widespread. Much of our talk about abstract entities – talk of models, possible worlds, functions, properties – is non-face-value, and subject to similar arguments as number talk.

3.2.2. Arguments from Linguistic Evidence

A popular strategy is to argue that the target discourse behaves in the same way as is typical of a particular type of non-face-value discourse (e.g. metaphor, pretence, myth, story-telling, and so on) and that this gives us positive reason to think that the target discourse belongs in that category of non-face-value discourse⁴⁴. Yablo (2001) argues that the best arguments for fictionalism appeal

⁴⁴ Note that this does not stand in tension with criterion iv. Here, the point is that a discourse domain D behaving like a fictional discourse F gives us reason to think that D is an instance of F. The notion of

to what features of the target discourse fictionalism is able to explain. Here is how he puts the point:

Just maybe, it gives the most plausible account of the practice. It is not that X's are intolerable, but that when we examine X-language in a calm and unprejudiced way, it turns out to have a whole lot in common with language that is fictional on its face. (Yablo 2001:87)

Unlike the ontological arguments discussed above, this strategy of argument looks at the discourse itself, and asks whether features of the discourse are best explained by positing a fictional mechanism. For example, Yablo (2010) argues that abstract entity discourse (including mathematical discourse) behaves like metaphorical discourse in many ways. Here are some of these ways, the first of which has already been touched upon:

Commitment

Metaphorical sentences do not commit speakers to their face-value truth, nor to any proposition or entity entailed by the truth of a face-value reading. Uttering “Juliet is the sun” does not commit Romeo to any propositions entailed by a face-value reading of this sentence, such as the proposition that Juliet is a blazing ball of plasma. Furthermore, uttering “Juliet is a bright angel” does not even commit Romeo to the existence of angels. In the same way, uttering “there were 12 apostles” does not commit speakers of mathematical discourse to the entailed claim that there are abstract platonic entities called numbers.

Impatience

In an example adapted from Burgess and Rosen (1997), Yablo (2010:148) asks us to imagine that the Oracle of Philosophy tells us that there are no abstract entities. Would we demand that mathematicians cease conducting their work, and that laypeople cease using number-talk? Presumably not. And if we did, we should expect to be met with impatience. This feeling of impatience is a datum to be explained. And indeed, speakers who utter metaphorical sentences naturally experience a certain feeling of impatience when otherwise competent interlocutors misread their utterances as communicating their face-value content. For example, a speaker who utters the sentence “their relationship is on the rocks” will naturally experience impatience if, barring some special explanation, an otherwise competent interlocutor consistently misreads the sentence as communicating a face-value content (“That can’t be right, there are no rocks around where they live”).

Silly Questions

Relatedly, metaphors stretch only so far. For this reason, they invite “outrageously inappropriate questions” (Yablo 2010:140) such as: exactly where is the hatchet buried? What virus did you catch to make you sick of all these roadworks? “Do you plan to *drop*-forge the uncreated

fiction would then be applied to explain how D fulfils its function. Criterion iv rules out views on which mere similarity between D and F is thought to suffice to make a view count as fictionalist.

conscience of your race in the smithy of your soul, or use some alternative method?” (Yablo 140). These questions are silly because they miss the sense in which metaphors are only as detailed as they have to be in order to illuminate the relevant target domain. In this sense metaphors are skeletal: there are no facts to discover about the hatchet, being sick of it, or even forging uncreated consciousness, beyond those minimally required to structure the target domain. As Yablo puts it, “why should we have filled out the story further than needed?” (2010:195). Yablo argues that number-talk is like this also. Questions such as “what are the intrinsic properties of the empty set? Are universals wholly present in each of their instances? Do relations lead a divided existence, parcelled out among their relata?” (2010:173) take us beyond the limit to which number-talk stretches⁴⁵.

That a domain of discourse demonstrates characteristics typical of a specific type of non-face-value discourse, then, is well explained by positing that the discourse is in fact an instance of that type of non-face-value discourse. For example, that abstract platonic entity talk such as number talk behaves like metaphorical talk is well explained by positing that such talk is in fact typically metaphorical.

3.4. Chung’s Epistemic Fictionalism

Returning to epistemology, Julianne Chung (2018, 2020) has recently argued that knowledge talk presents us with a puzzle similar to the puzzle presented by mathematical talk, and that just as in the mathematical case, being fictionalists about knowledge resolves the problem. To be a fictionalist about knowledge talk is to hold that positive sentences of the form “S knows that P” do not communicate their face-value meaning, but aim to achieve some alternative aim. Chung’s version of fictionalism is an *assertoric fictionalist* view (Chung uses the term ‘content fictionalism’ to describe her view). It holds that the alternative aim achieved by uttering sentences of the form “S knows that P” is to assert some non-face-value content, a content which differs from the face-value content. Crucially, the non-face-value content is not only different to the face-value content, but it might actually be consistent with the contradiction of the face-value content.

Here is how Chung (2018:391) characterises her view:

[S]entences of the form “S knows P” (and the like) usually, if not always, communicate false propositions. Does this mean that we are not warranted in our use of such sentences? No, says the Fictionalist. For, we do not normally use a sentence of the form “S knows P” to assert the proposition that it semantically communicates. Rather, we use it non-literally in such a way so as to avoid asserting its semantic content and to assert some other content (consistent with the proposition that S does not know p) altogether.

⁴⁵ A real life example of this phenomenon: as I write this chapter my son, who has just had his third birthday and therefore is no longer two years old, wants to know “where ‘the two’ went”. These questions are not silly when uttered by young children precisely because young children have not yet developed the ability to “see through” the metaphor to the communicated non-face-value content.

So, to be a fictionalist about knowledge, one divorces the face-value content (the meaning lodged into surface grammatical form and semantic content) from non-face-value content (the meaning that such sentences are used to actually communicate).

On Chung's view, the non-face-value content communicated by a typical utterance of sentences of the form "S knows that P" is deeply context-sensitive, and varies with use:

Typically asserted propositions may include those along the lines of: that S truly believes P, that S is justified in believing P, that S believes P as a result of a reliable belief-forming process, that S can rule out all relevant alternatives to P, that S can assert P, that S can rely on P in practical reasoning, and so on – in other words, propositions that many epistemologists have argued are associated with knowledge-attributing sentences. However, it should be emphasized that these may not even come close to exhausting the propositions that we use sentences of the form "S knows P" to assert. (2018:392)

In contrast to this, the face-value content is the one posited by the invariantist. Chung takes the face-value meaning to be "the proposition that S knows that P", where 'knows' is given an invariantist treatment. However, she leaves open the question of just how high the standards for knowledge are on the face-value reading (2018:405):

[i]t seems as if "knows" quite possibly does have a sober – and yet highly, if not impossibly, demanding – use in a significant number of contexts.

Continuing in a footnote (2018:405):

Indeed, perhaps so highly demanding that it is impossible to meet the satisfaction conditions that it mandates (which arguably rise to the level of infallibility).

One way in which sceptical argument might render face-value readings infelicitous is by being sound, thereby rendering knowledge sentences false if taken at face value. However, this is not the only way. As Chung points out, the fictionalist need not commit herself to the view that sceptical arguments are in fact sound if knowledge sentences are taken at face-value. It suffices that their being sound is a serious possibility and that whether they are sound is unclear from our perspective. From this it would follow that whether our knowledge sentences are false is also unclear, which is enough to make knowledge sentences infelicitous on a face-value reading. Put differently, it would suffice for the standards of knowledge to be high, in a way that nevertheless falls short of infallibility, provided that it were obscure to us how they fell short of infallibility.

So on Chung's view, sentences of the form "S knows that P" typically communicate the face-value content that S knows that P, in order to communicate the non-face-value content, which takes the form of a disjunction and that I will call "Justified True Belief + (JTB+)".

JTB+ epistemic fictionalism

Face-Value content:	S knows that P [invariant, possibly sceptical infallibilism]
Non-Face-Value content:	JTB+ [S truly believes P, that S is justified in believing P, that S believes P as a result of a reliable belief-forming process, that S can rule out all relevant alternatives to P, that S can assert P, that S can rely on P in practical reasoning, and so on]

Chung's fictionalist view is intended to motivate fictionalism as a general approach, and leaves open the question of what underlying fictionalist mechanism is underpinning the redirection from face-value content to non-face-value content. Thus, we can see Chung as providing a road map and initial motivation for a fictionalist view, to be filled in and fleshed out with a specific fictionalist mechanics.

Chung motivates her view primarily by appeal to linguistic evidence, arguing that knowledge talk exhibits some of the hallmarks of non-face-value talk outlined by Yablo. In particular, she argues that knowledge talk too appears to leave a gap between what speakers intend to commit themselves to and what (a face-value reading of) their uttered sentences seems on the face of it to commit them to, as evidenced by the urge to retract that speakers feel upon being presented with sceptical arguments. Furthermore, she argues that sceptical hypotheses are examples of questions that appear legitimate but which nevertheless elicit impatience and strike speakers as silly. Just as is the case for mathematical talk, we expect that everyday users of knowledge talk to feel impatient if we started to insist that they cease making knowledge attributions until the problem of scepticism has been resolved; and in many everyday contexts the question "are you sure you're not a brain in a vat?" strikes us as an unserious response to a good faith knowledge attribution.

3.5. Applying the 5 Desiderata

Let's return to the five desiderata for a well-motivated fictionalist theory. How does epistemic fictionalism fare? We have already seen that epistemic fictionalism meets both criteria i. and ii. epistemic fictionalism is (i.) a theory about fragments of discourse, specifically, knowledge talk; and (ii.) holds that due to the possible soundness of sceptical arguments, knowledge discourse is semantically infelicitous on a face-value reading. Chung does not address (iii.) and (iv.) directly, but it is easy to fill in the gaps. On Chung's account, knowledge sentences communicate propositions such as 'S has a justified belief that P'. Given the overwhelmingly plausible hypothesis that evaluating beliefs epistemically is useful, redirecting towards this non-face-value content explains the function of knowledge discourse. Chung's epistemic fictionalism, then, meets criteria (i.)-(iv.).

However, I think that epistemic fictionalism as it has been formulated fails on criterion (v.). One reason is that Chung's fictionalism is intended as a sketch to be filled in with a relevant mechanism, and so does not specify the notion of fiction in play. Therefore, it naturally also does not explain how a particular notion of fiction might explain the posited redirection. This is fine, given the 'roadmap' aim of the paper. The deeper problem is that epistemic fictionalism needs

to account for why sentences of the form “S knows that P” would be non-face-value – why communicate that one has, for example, a justified belief that P in terms of a fiction in which one is infallible with regards to p? What is the advantage of using the metaphor of infallibilism to communicate that one is in some other, achievable state? Why not simply say that S has a justified belief that P, where ‘S has a justified belief that P’ is taken at face-value, and forego non-face-value talk of knowledge altogether? In this section, I argue that an epistemic fictionalism that posits JTB+ as its communicated content cannot answer these questions. None of the suggested functions – representational, presentational, or procedural- are applicable to knowledge sentences if knowledge sentences communicate JTB+ on a non-face-value reading.

3.5.1. Representationally essential discourse

Representationally essential discourse functions to communicate content for which there is no face-value paraphrase. Literary metaphors are often of this genre: fiction is employed to fill a void that no exclusively face-value proposition communicates. As Yablo (2010:135) puts it: “the language might have no more to offer in the way of a unifying principle for the worlds in a given content than that *they* are the ones making the relevant sentence fictional. It seems at least an open question, for example, whether the clouds we call angry are the ones that are literally F, for any F other than 'such that it would be natural and proper to regard them as angry if one were going to attribute emotions to clouds'”. Other examples include piercing glances, warm laughs, tangled lines, heartbreak, cutting words. Perhaps there is no face-value proposition that communicates what it is to be an angry cloud, a piercing glance or, in the case that interests us, knowledge that P. Toon’s (2016) psychological fictionalism is an example of a fictionalist theory that claims to concern representationally essential discourse. Toon posits that talk of belief and desires is not to be taken at face-value and functions to predict and explain agents’ overt behaviour. Yet, argues Toon, there is no specific feature shared by the wide range of overt behaviour characterised in terms of belief and desire, other than being understood and predicted in terms of belief and desire.

Knowledge discourse is not representationally essential because there are face-value paraphrases for knowledge, and these paraphrases are the content that knowledge communicates (JTB+). The non-face-value term ‘knowledge’ therefore does not exist to represent something that can only be communicated using non-face-value content. Note that the possibility of the propositions that knowledge communicates being elusive and inaccessible does not suggest that knowledge is representationally essential. A proposition being elusive means that speakers are not always sure *which* non-face-value content a knowledge sentence communicates, not that there is no non-face-value content. Further, sentences in a representationally essential discourse could not even in principle be communicated in a face-value way. If knowledge too could not even in principle be communicated in terms of non-face-value propositions, then the claim that knowledge is context sensitive and communicates different non-face-value propositions depending on the context would be incoherent.

3.5.2. *Procedurally essential discourse*

Procedurally essential non-face-value discourse functions to enable agents to think and communicate about important information in cases where the ontology underpinning the information is unimportant or unobvious. Here is Yablo (1998:168) on procedurally essential talk:

[O]ne of the contents that my utterance may be up for, when I launch S into the world in the opportunistic spirit described above, is its literal content. I want to be understood as meaning what I literally say if my statement is literally true (count me a player of the ‘null game’, if you like) and meaning whatever my statement projects onto via the right sort of ‘non-null’ game if my statement is literally false. It is thus indeterminate from my point of view whether I am advancing S’s literal content or not.

Isn’t this in fact our common condition? When people say that the number of apostles is twelve [...] they are far more certain that S is getting at something right than that the thing it is getting at is the proposition that S, as some literalist might construe it. If numbers exist, then yes, we are content to regard ourselves as having spoken literally. If not, then the claim was that there were twelve apostles.

Procedurally essential non-face-value discourse goes unnoticed precisely because whether the utterance is true if taken at face-value is indeterminate from the speaker’s perspective. Perhaps such things as platonic numbers exist – the speaker certainly is not committed to it – in which case the speaker intends the utterance to be understood at face-value. If not, the speaker intends to be understood as communicating the sentence’s non-face-value content. The speaker uses number talk to aid her practical tasks, but the speaker holds no particular commitments regarding the truth of the face-value reading. Thus, it is largely unimportant to speakers which content they are communicating. Although he does not couch it in these terms, Le Poidevin’s (2020) religious fictionalism contains an example of procedurally essential discourse. For Le Poidevin, fictionalism enables the agnostic to engage in religious life while remaining neutral about whether the tenets of her religion are true at face-value or whether they are useful fictions.

However, knowledge talk does not appear to fit this pattern. If knowledge talk were procedurally essential then we should expect speakers to be non-committal regarding whether the face-value reading of sentences of the form “S knows that P” is true. According to the mathematical fictionalist and the moral fictionalist, speakers who use number-talk and moral-talk are non-committal about the existence of abstract platonic entities and attitude independent moral facts. From the perspective of such speakers, it is indeterminate whether these entities exist: if they do exist, then speakers are happy to take their sentences to commit them to the existence of such entities, but if they do not then speakers do not wish to be thusly committed. But this is not plausible in the case of knowledge talk: in the case of knowledge talk, we do *not* wish to remain non-committal about the face-value truth of sentences of the form “S knows that P”. Rather, we

wish to disavow commitment to being infallible with regards to the propositions that we believe⁴⁶. Furthermore, if sentences of the form “S knows that P” communicate something like “either S is infallible with regards to P or S is JTB+ with regards to P” the project in epistemology of resisting scepticism would be at odds with everyday speaker meaning. If sceptical acceptance is built into what everyday knowledge talk is intended to express, then it would be mysterious why we should ever have tried to resist sceptical arguments.

3.5.3. *Presentationally essential discourse*

Presentationally essential non-face-value talk enables us to think and talk about propositions that would otherwise be difficult or perhaps even impossible to think and talk about. This is achieved by presenting the target propositions in non-face-value ways that are more cognitively and psychologically accessible than the non-face-value paraphrases would be. Take for example the sentence ‘the average star has 2.4 planets’. There is no such thing as an ‘average star’ – such an entity is a fiction – but thinking about the relationship between the number of stars and the number of planets in terms of averages is less cognitively demanding and more psychologically accessible than thinking about the long disjunction that would result if we paraphrased talk of averages away. As Yablo (1998:253) puts it, “to think of Crotona’s location as the place it would need to be put in the arch of Italy imagined as a boot, or as the stars and planets to be proportioned the way they would need to be for the average star to come out with 2.4 planets, is to be affected in ways that go way beyond the propositions expressed”.

That a non-face-value domain of discourse is presentationally essential is a popular function posited for the targets of fictionalism. For example, Levy’s (2011) fictionalist account of information and communication discourse in biology holds that genes, hormones, morphogens and other entities are often described as transmitting and receiving signals in order to present information about these entities’ causal roles within wider biological systems in more accessible ways. Eshleman’s (2005) religious fictionalism can also be understood as positing religious discourse as presentationally essential. Eshleman argues that presenting the ideal of a perfect state of being via the non-face-value concept of God motivates us to improve ourselves by framing a striving towards perfection.

Applying this model to Chung’s epistemic fictionalism would be to claim that presenting JTB+ via the language of infallibilism enables us to better think and reason about JTB+. Why think, however, that presenting JTB+ as knowledge better positions us to reason about JTB+? On the contrary, it seems that JTB+ is *easier* to reason about than the term ‘knowledge’. Talk of ‘knowledge’ makes it less clear whether a particular belief is an instance of JTB+, compared with using the relevant face-value proposition. As evidence, consider the diverging intuitions elicited by certain Gettier cases, such as whether Henry knows that he is looking at a barn when he’s in

⁴⁶ This might seem too quick. Unger (2002) and Dodd (2011) interpret everyday speakers as thinking that there is a possibility that they infallibly know many things. However, if everyday speakers do take themselves to be infallible with regards to ordinary propositions, then fictionalism is redundant. On this view knowledge sentences would not be semantically infelicitous on a face-value reading, and the second fictionalist criterion would not be met. Fictionalism, then, requires that people do not take themselves to know infallibly.

fake barn county (Gettier, 1963). Crucially, the divergence concerns Henry's knowledge: epistemologists and everyday speakers alike tend to agree that Henry has a justified true belief that he is looking at a barn. On the assumption that epistemologists and everyday speakers are correct in their justification judgements, that JTB+ intuitions converge while knowledge intuitions diverge suggests that reasoning about JTB+ is easier than reasoning about knowledge. Further, unlike talk of average stars and Italian coastal town locations which require complex calculations and detailed familiarity with Italian geography to communicate in face-value terms, reasoning about JTB+ does not require particular ability or specialist knowledge that might need to be presented to better advantage with the concept of knowledge.

The second reason to doubt that knowledge function to present JTB+ in more cognitively advantageous ways is that there are counter-examples to JTB+ that are not also counter-examples to non-face-value readings of knowledge. Consider another Gettier case: Harriet sees a sheep-shaped rock in a field, and forms the belief that there is a sheep in the field. In this case, there actually is a sheep behind the sheep-shaped rock, so her belief is justified and true - but nevertheless, fails to be knowledge (Chisholm, 1966). The problem is that the Sheep Case is incompatible with knowledge but compatible with at least two of the contents in JTB+, even on a non-face-value reading of knowledge. The face-value reading of knowledge is that posited by the infallibilist, and so we should expect many counter-examples to knowledge that are not counter-example to JTB+ on a face-value reading of knowledge⁴⁷. The problem for JTB+ epistemic fictionalism is this: even when knowledge is being used in a non-face-value way it seems as though the Sheep Case fails to be knowledge, *without* failing to be instances of at least some of the contents in JTB+. Compare this to a non-controversial metaphor case. Say that 'bury the hatchet' is typically used in a non-face-value way to communicate that a dispute has ended. We should expect many instances of ending a dispute that are not instances of burying the hatchet, where burying the hatchet is given the face-value reading of 'to put an axe in a hole in the ground and cover it with dirt'. However, there should be no instances of ending a dispute that are not also instances of burying the hatchet on a non-face-value reading. If 'burying the hatchet' functions as a suggestive way of communicating that a dispute has ended then any instance accurately described by 'the dispute ended' is also accurately described by a non-face-value reading of 'bury the hatchet'. Likewise, if "S knows that P" functions as a suggestive way of communicating that 'S JTB+s that P' then any instance accurately described by the propositions in JTB+ will be accurately described by a non-face-value reading of "S knows that P".

⁴⁷ As touched on above, metaphors typically operate by 'partial mapping'. This means that some key elements of the source domain are mapped onto the target domain, while other elements are not. For example, the metaphorical mapping 'recovery is a journey' maps some of the elements of the concept 'journey' onto the concept 'recovery' (that being cured is a 'destinations' towards which the agent is moving; that it is possible to take a 'wrong turn'; that the people met in the context of recovery are 'journey companions') while other elements of the concept journey are not mapped onto the concept 'recovery' (that driving is faster than walking; that it is prudent to check ones tyres before leaving; that one should remember to pack colouring books with which to entertain young children). We should, therefore, expect there to be examples that are true of the non-mapped elements of the source domain, but false of the target domain, and vice versa.

This section has argued against the view that knowledge communicates JTB+ by showing how knowledge as a non-face-value way of communicating JTB+ does not fulfil any of Yablo's suggested functions for fiction. Given the widespread popularity of these functions, this goes some way towards showing that knowledge as a non-face-value way of communicating JTB+ does not fulfil any functions, and therefore that JTB+ is not the non-face-value content typically communicated by knowledge sentences. These arguments do not show that there are no other, as of yet unexplored functions that knowledge as a non-face-value way of communicating JTB+ might fulfil. I do however think that the prospect of such unexplored possibilities is pessimistic. If the above arguments are correct, then reasoning about JTB+ is no more onerous than reasoning about knowledge. The next chapter will defend an alternative communicated content for a non-face-value reading of sentences of the form "S knows that P", and argue that this alternative content makes sense of the function of knowledge talk.

3.6. A Refinement: Chung's Conceptual Metaphor Theory Epistemic Fictionalism

In 'Skeptical Arguments, Conceptual Metaphors, and Cross-Cultural Challenges', Chung (2020) explores a refinement to the general roadmap outlined in her 2018 paper. In this more recent work she specifies the mechanism by which redirection takes place as involving metaphor, and in particular, metaphor understood along the lines of Lakoff and Johnson's Conceptual Metaphor Theory.

According to Conceptual Metaphor Theory (hereafter, CMT) the vast majority of our discourse is metaphorical. Lakoff and Johnson understand metaphor as consisting of pairs of concepts: the *topic*, which corresponds to the feature of the world that the metaphor applies to ("Juliet" in "Juliet is the sun") and the *vehicle*, which corresponds to the concept that augments how we experience the topic. Crucially, on CMT, metaphorical talk reflects metaphorical thought and experience. In everyday metaphors, the topic is typically experienced as the vehicle: "[o]ur ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature" (Lakoff and Johnson, 1980: 3).

To illustrate with Lakoff and Johnson's own leading example, take the target concept of an *argument* and the conceptual metaphor ARGUMENT IS WAR. The idea is that everyday expressions such as "attack his argument", "win an argument", and "defend a key claim" communicate an underlying relationship between the concepts argument and war, that can be summarised in the conceptual metaphor ARGUMENT IS WAR. This is not to say that arguments are subspecies of wars, but rather that what it means to be an argument is to be an activity that is structured along the lines of war. These metaphors are not merely linguistic flourishes, but reflect the fact that we experience arguments as, in some sense, wars. To give another classic example, "She has a warm smile" reflects the underlying conceptual metaphor AFFECTION IS WARMTH, whereby we experience emotion and affection as in some way connected to (physical) warmth. Other reflections of this metaphor include its reverse: "Sally is a block of ice" and "a chilly reception".

CMT is at its heart an embodied theory of meaning. Almost all of our concepts are metaphorical on this view. Lakoff and Johnson (1980:3) write that because "our conceptual system is largely

metaphorical, then the way we think, what we experience, and what we do every day is very much a matter of metaphor”. Those concepts that are not themselves metaphorical are “simple spatial concepts” (1980:56) such as UP, DOWN, WARM, COLD, LIGHT, DARK, and so on. In short, simple concepts are embodied concepts, and these are what ground metaphorical, non-embodied concepts. The way that we experience the world from childhood into adulthood lays the foundations for the conceptual metaphors that structure our experience. The experience of being held and comforted by a loving parent, for example, lays the foundations for the AFFECTION IS WARMTH conceptual metaphor that goes on to structure how one experienced affection, even into adulthood.

Returning to epistemic fictionalism, then, Chung argues that knowledge thought and talk is experienced along the lines of the following conceptual metaphor: BELIEVING IS KNOWING, where ‘knowing’ is cashed out in terms of psychological certainty. The key claim is that some types of believing are experienced as certainty, where this is marked by a psychological feeling of certainty; and these are the beliefs that are counted, metaphorically, as knowledge.

CMT epistemic fictionalism

Vehicle:	S is (psychologically) certain that P
Topic:	S believes that P is true

Chung argues that as was the case for AFFECTION IS WARMTH, “experiences of thinking-true—or (for our purposes here) believing—are often accompanied by a sensation of certainty: or, what we might call a literal feeling of knowing” (2020:170). Here, sensations of psychological certainty and feelings of knowing just are the same thing. During childhood the feeling of at least certain kinds of believing become associated with psychological certainty, building up an association between the domain of BELIEVING and the domain of KNOWING (understood, at face-value, as psychological certainty). This is not to say that believing is a subspecies of knowing (understood as psychological certainty) but rather that even once we learn to differentiate between believing and feeling psychologically certain, that experiencing belief is somehow structured along the lines of experiencing certainty. The strong phenomenological association between believing and feeling certain leads us to characterise (at least some) instances of believing as instances of knowing.

Note that although this picture adds a particular fictional mechanism (vehicle and topic, as conceived by CMT), it retains the picture of communicated content of the original view. On this picture, too, sentences of the form “S knows that P” communicate the contents JTB+: that S truly believes P, that S is justified in believing P, that S believes P as a result of a reliable belief-forming process, that S can rule out all relevant alternatives to P, that S can assert P, that S can rely on P in practical reasoning, and so on.

3.6.1. Objections

I will discuss three brief objections to this refinement of the view, before moving on to the version of fictionalism which I prefer.

The first point to note is that on this view, the posited communicated content is the same as the original view (JTB+); and so all of the objections regarding the function of the communicated content apply to this new view also. What's more, one might worry that this way of understanding knowledge talk actually makes knowledge talk *less* functional. For psychological certainty is not always a good guide to whether a belief is likely to be justified, reliable, and so on. If the concept of knowledge really is structured primarily by mapping feelings of psychological certainty onto beliefs, then at its core knowledge talk is structured on how we feel about our beliefs, as opposed to being structured in a way that relates to what epistemic position we might be in. Indeed, Chung recognises that this concept of knowledge might not be as functional as we'd hope, writing that perhaps "we just cannot help ourselves from conceptualising our beliefs in this way (given the conceptual metaphor's basis), even if we recognise that doing so is potentially misleading and even dangerous in certain kinds of ways" (2020:171). Although it fits into the wider CMT perspective, I think that positing a poorly functioning concept of knowledge speaks against a fictionalist theory, as per criteria (v).

It should be noted that Chung (2020:171) does also make some suggestive comments regarding positive functions of metaphorically mapping believing in terms of certainty, writing that:

Using "knows" and "knowledge" metaphorically also may encourage us to be curiously skeptical, rather than paralyzed by doubt; as such it may permit us to act as if we are immune from error (at least to an extent, just as conceiving of arguments in terms of war permits us to act as if arguments are war, at least to an extent), thereby allowing us to avoid the pitfalls of excessively entertaining doubts while at the same time shielding us from the dangers of dogmatism.

I think that these comments are along the right lines, and will return to these ideas in the next section. However, conceiving belief as mapped in terms of psychological certainty in particular makes it difficult to see why we would be shielded from the dangers of dogma, and how we might apply the mapping so that only the right types of beliefs are mapped in terms of certainty.

A second problem is that (without further explanation) it is unclear how and why mapping feelings of knowing onto belief succeeds in generating the communicated contents captured by JTB+. For a start, it is unclear why the feelings of certainty would apply only to those beliefs that are justified, reliable, and the like. Perhaps there is a story to be told on which as time goes on, feelings of certainty become more attuned to these good-making features of particular beliefs, and we come to experience psychological certainty primarily when epistemic certainty is warranted⁴⁸. But this still would not go the full way towards explaining how the mapping generates the posited content in the first place: specifically what is it about mapping certainty onto belief that redirects sentences of the form "S knows that P" so that they come to communicate "SJTB+s that P"? This problem is particularly salient when focusing on third person knowledge attributions. Even if it is the case that the BELIEVING IS KNOWING mapping structures sentences

⁴⁸ To be clear, I do think that emotional valence is one of the hallmarks of metaphorical content (something that I return to in Chapter 4). But I think that mapping knowledge solely in terms of certainty generates the content JTB+.

of the form “I know that P” so that we map feelings of psychological certainty onto certain types of beliefs, it is somewhat mysterious how the mapping structures sentences of the form “S knows that P”. For whether S is psychologically certain that P is not something that we would expect the speaker to readily have access to. Additionally, we would not expect whether S is certain as to whether P to be clearer and more easily accessible to us than whether S believes that P. Again, perhaps there is a story to be told about how we move from our own feelings of certainty to make inferences about those of others – but this story begins to move away from the embodied story on which speakers use their feelings to structure how they experience believing. Now it looks like speakers are working out whether S does (or should) feel certain as to whether P, which looks like a more complicated and less embodied story. Notice how, in contrast, the AFFECTION IS WARMTH mapping is far more straightforward to apply: it concerns solely how the person or action being described as warm or cold makes the speaker feel.

Finally, I wish to flag a key commitment incurred by any fictionalist view based on CMT: on this view, all of our philosophical concepts are inherently metaphorical. CMT posits that there are two types of concept: non-metaphorical concrete physical ones (e.g. *up, down, empty, forward*) and metaphorically structured abstract ones (e.g. *love, happiness, time, money, freedom, equality*). Here is the catch: on CMT, almost all of our concepts are abstract, and therefore metaphorical. Certainly all our philosophical concepts are metaphorical. Indeed, Johnson (2012:39) lists knowledge as an “irreducibly metaphoric” concept alongside the concepts *causation, events, will, thought, reason, mind, justice, and rights*. Noting this commitment isn’t in itself an objection either to CMT or to Chung’s fictionalism. But there are two upshots to note: first, Chung’s way of making good on the claim that knowledge talk is metaphorical commits her additionally to the view that all philosophical concepts, and many lay concepts too, are metaphorical. If you buy into CMT for knowledge talk, you need to also buy into it for other philosophical concepts, or else the explanation for why knowledge talk is metaphorical loses its motivation. Second, if CMT is the right account of how language and thought works in general, then knowledge talk is, trivially, metaphorical as well. The claim that knowledge too is metaphorical is not in itself a radical proposal, but one that falls out of accepting CMT as a general framework. In the next chapter, I will defend a view of fictionalism that does not rely on CMT. I take the fact that my view does not commit me to the claim that all philosophical concepts are metaphorical to be an advantage of my way of construing epistemic fictionalism.

3.7. Conclusion

This chapter has introduced fictionalism in philosophy, and defended some characterisations. I have tried to show how function-first is a natural method for a motivated fictionalist theory: looking at the Fictionalist’s target discourse with an eye trained on the function of the discourse, and on how the discourse being fictional enables it to fulfil this function, aids the fictionalist’s case for thinking that discourse involves fiction. I then argued that Chung’s version of epistemic fictionalism, on both formulations, does not explain why we should expect to have a fictional concept of knowledge, such that sentences of the form “S knows that P” come to communicate the non-face-value content JTB+. The next chapter develops an alternative epistemic fictionalist picture, and argues that it succeeds in explaining why the posited content would be communicated via fiction.

Chapter 4: Epistemic Fictionalism, Function-First

This chapter develops a way of being a fictionalist about knowledge that centres the purpose of the concept of knowledge. This view, which I call Function-First Epistemic Fictionalism, holds that sentences of the form “S knows that P” express the face-value content posited by the sceptical infallibilist: perhaps, that S has ruled out every epistemic possibility in which $\sim P$, that S’s evidence E entails that P, that P has a probability of 1 on S’s evidence, and so on. Yet, such sentences are typically used to communicate another, non-face-value content, which is directly related to the roles that knowledge plays: perhaps, that S is a reliable informant with regards to P, S may appropriately terminate inquiry into whether P, S ought to be honoured for having achieved her particular epistemic position with regards to P, S is blameworthy if she acts on P to the detriment of others, and so on. In particular, I argue that knowledge talk is metaphorically framed, so that speakers think and talk about the communicated content in terms of the face-value content.

The chapter starts by outlining my novel function-first epistemic fictionalist proposal (§4.1.). I argue that function-first epistemic fictionalism is to be preferred to Chung’s fictionalist account, on the basis that it meets all five of the criteria for a well-motivated fictionalism outlined in the last chapter, while retaining the advantages of Chung’s original view (§4.2.). Then, I move on to one of the key contributions of this thesis: putting epistemic fictionalism to work in resolving the problem presented by sceptical infallibilism (§4.3.). I argue that the face-value reading of sentences of the form “S knows that P” is that posited by the sceptical infallibilist, which explains why sceptical infallibilism is compelling: it is, literally, true. Yet, epistemic fictionalism avoids the problems commonly associated with sceptical infallibilism: sentences of the form “S knows that P” typically communicate a true content, and so everyday speakers are not committed to widespread error. Finally, I explain how epistemic fictionalism retains the advantages of sceptical infallibilism outlined in Chapter 1 (§4.4.).

4.1. Function-First Epistemic Fictionalism: a proposal

Before spelling out this function-first fictionalist proposal, here is a re-cap of the features of function-first epistemology that will play a central role in supporting the proposal. Function-first epistemology (a) seeks to elucidate the concept of knowledge by investigating the social functions that it fulfils; (b) derives an analysis of knowledge by reflecting on the intension that knowledge would require in order to successfully fulfil the posited functions; and (c) often concludes that the standards for knowledge are sufficiently demanding such that they enable knowledge to fulfil its social roles without being so demanding that knowledge is unobtainable (Craig 1990, Fricker 2008, Hannon 2019a, 2019b). Say that knowledge serves the role of flagging which members of the community are reliable informants with regards to a particular proposition (Craig 1990, Hannon 2019a, Pritchard 2012). The standards for flagging a potential informant better be demanding enough to prevent unreliable potential informants from being flagged as reliable, without being so demanding as to exclude informants who, while fallible, are reliable enough for most ordinary purposes. Similar considerations apply to any of the posited functions for knowledge. Contra (c), Chapter 2 argued that knowledge’s functions could still be fulfilled if infallibilism is true. This chapter proposes a middle ground on which sceptical infallibilism

captures the face-value content of knowledge sentences, while non-sceptical, function-driven fallibilism captures the non-face-value content that knowledge sentences typically communicate.

4.1.2. *Non-Face-Value and Face-Value Content*

This fictionalist view, which I call Function-First Epistemic Fictionalism, shares with Chung's fictionalism the idea that if taken at face-value, sentences of the form "S knows that P" are rarely, if ever, true. The face-value content is that posited by the sceptical infallibilist: that S has ruled out every epistemic possibility in which $\sim P$, that S's evidence E entails that P, that P has a probability of 1 on S's evidence, and so on. Yet my account differs from Chung's in an important respect. We have seen that for Chung, sentences of the form "S knows that P" typically communicate JTB+ (that S truly believes P, that S is justified in believing P, that S believes P as a result of a reliable belief-forming process, that S can rule out all relevant alternatives to P, and so on). In contrast, I want to suggest that the non-face-value contents typically communicated by knowledge sentences directly reflect the functions of knowledge talk: that S is a reliable informant with regards to P (Craig 1990, Hannon 2019a, Prichard 2012); that S may appropriately terminate inquiry into whether P (Kappel 2010, Kelp 2011, Rysiew 2012); that S ought to be honoured for having achieved her particular epistemic position with regards to P (Kusch 2009); that S is blameworthy if she acts on P to the detriment of others (Beebe 2012); and so on. In other words, the propositions that many function-first epistemologists posit as capturing the purpose for knowledge talk. Call these propositions taken together *Reliable Informant +* (RI+).

This is not to suggest that these propositions represent every possible content in RI+, and it certainly is not to suggest that utterances of sentences of the form "S knows that P" communicate *all* of these contents simultaneously. Rather, the fictionalist framework is compatible with a range of accounts of the function(s) of knowledge, and hence, compatible with a range of different potential communicated contents. Many function-first epistemologists think that knowledge serves more than one function (e.g. Beebe 2012, McGrath 2015, Hannon 2019a, Kusch and McKenna 2020). If this is right, then the function-first epistemic fictionalist account preserves the context sensitivity posited in Chung's account. On this view, the content communicated by "S knows that P" varies depending on the context, sometimes communicating that S is a reliable informant with regards to P and sometimes communicating that S may end inquiry into whether P. For ease of exposition, I primarily focus on the flagging reliable informants function, and talk as if sentences of the form "S knows that P" communicate that "S is a reliable informant with regards to P". However, the view is compatible with alternative functions, including function pluralism.

In a nutshell, then, function-first epistemic fictionalism says that sentences of the form "S knows that P" would, if taken at face-value, communicate that "S is infallible with regards to P". However, they are instead typically used to communicate the non-face-value content that "S *RI+*s that P".

Function-First Epistemic Fictionalism

Face-Value content: **Infalibilism** [S has ruled out every epistemic possibility in which $\sim P$, S's evidence for P entails that P, P has a probability of 1 for S, S is epistemically certain that P, and so on.]

Non-Face-Value content: **RI+** [S is a reliable informant with regards to P, S may appropriately terminate inquiry into whether P, S ought to be honoured for having achieved her particular epistemic position with regards to P, S is blameworthy if she acts on P to the detriment of others, and so on.]

On this picture, sentences of the form "S knows that P" have at least two different readings: a face-value reading and a non-face-value reading. Whether such sentences are true depends at least in part on what reading is under discussion: a sentence may be true on its non-face-value reading, while being false on its face-value reading.

An initial worry is that we might wonder why the non-face-value content should communicate a content that *corresponds* to the function, as opposed to communicating a content that is *shaped* by the function. Function-first epistemologists generally hold that keeping an eye on the function served by the concept of knowledge helps us to work backwards to illuminate what knowledge must be like. But claiming that it is illuminating to analyse knowledge by thinking about how it enables us to recommend reliable informants is very different from claiming that "S knows that P" *means* that "S is a reliable informant as to whether P". Rather, function-first epistemologists typically use their function regarding the function of knowledge to argue that knowledge must have a certain shape. That is to say, function-first philosophers typically think that sentences of the form "S knows that P" communicate the content that a given theory of knowledge might predict (e.g. that contextualism or invariantism might predict⁴⁹) rather than thinking that sentences of the form "S knows that P" directly communicate that S is a reliable informant as to whether P.

Yet function-first epistemologists asks: what can we do with this concept that we could not, or could not *so easily*, do without it? And (at least some) answer: we could not recommend reliable informants to one another. So if what we do with sentences of the form "S knows that P" is recommend informants, why not think that what we communicate when we communicate these sentences just is that S is a reliable informant? After all, if communicating this content is the purpose of the very concept of knowledge, why not think that the concept is shaped such that it does directly communicate this content? The main reasons that the realist function-first epistemologists cannot straightforwardly claim that knowledge talk means the content of RI+ is that this would make it difficult to explain the appeal of scepticism. For if sentences of the form "S knows that P" just mean "S is reliable as to whether P" then it would be difficult to see how

⁴⁹ For a function-first defence of contextualism, see Greco (2008), Henderson (2009, 2011), McKenna (2013) and Hannon (2013, 2015. Although cf. 2019a); see Grimm (2015) for a defence of sensitive invariantism and Kelp (2011) for a defence of insensitive invariantism.

the sceptical infallibilist could ever persuade us that we do not possess lots of knowledge: it seems apparent that many people do routinely qualify as reliable, if fallible, informants. Second, it just doesn't seem to users of knowledge talk that this is what they mean when they utter sentences of the form "S knows that P". As Queloz (2021) notes, functions are often hidden from view, and it takes philosophical investigation to uncover what they are. If the function of knowledge were straightforwardly communicated, then we might not expect the function to be invisible to speakers in this way. However, both of these objections are easy to dissolve on metaphor fictionalist views. First, scepticism targets the face-value meaning of sentences of the form "S knows that P", so (as I argue in §4.3.) we have an explanation for why sceptical worries get a hold, even though what knowledge sentences communicate is that S is a reliable informant with regards to P. Second, it is natural that speakers cannot always clearly identify what non-face-value content is communicated by an utterance. Part of what metaphor frames do is colour and structure how we think about the non-face-value content communicated, so that it is not always easy to think about the non-face-value content in terms of literal paraphrases. That is to say, it is not always easy, via introspection alone, to identify what precisely a non-face-value communicated content amounts to (a point I defend in Chapter 5).

4.1.3. *Metaphor Frames*

Having outlined the posited content, I now turn to the fictionalist mechanism which I take to underlie the redirection from the face-value to the non-face-value content. Following Chung (2020) who herself takes her cue from Yablo (1998), I think that the relevant fictionalist machinery at work in knowledge talk involves metaphor. However, I wish to take a more ecumenical approach than Chung does, and understand the process underpinning the redirection as involving *metaphor frames* (Lakoff 2001; Semino 2008, Cameron et al. 2010, Ritchie 2013, Camp 2006, 2007, 2019).

Metaphor frames influence the ways in which we think and talk about particular concepts and practices. In essence, metaphor frames involve a conceptual mapping between a source domain (the domain used in the metaphor) and the target domain (the domain being described metaphorically) such that key aspects of the source domain are used to make salient analogous aspects of the target domain. The idea is that framing one concept (the target domain) metaphorically in terms of another (the source domain) influences our thinking and augments our ability to reason about the target concept.

Take the well-trodden example of "Juliet is the sun". Here, Romeo applies a metaphoric framing between the target domain (Juliet) and a source domain (the sun) to give the audience to understand a metaphorical content. Precisely what this content is might not be easy to determine, but it might be something like: "that Juliet worthy of worship; that she is the focus of his thoughts and dreams; that she is an exemplar of goodness and beauty; that her goodness and beauty are natural and original, and far exceed those of the other ladies of Verona; and that her goodness nurtures Romeo and helps him to grow emotionally" Camp (2006b:3). Using the frame of the sun highlights and makes salient key features of Juliet: her warmth, her radiance, her centrality. In this way, Romeo exploits features of the source domain to get his hearers to think about Juliet in the way suggested by the source domain.

Literary metaphors are not the only place that we find metaphor frames. We can understand Lakoff and Johnson's leading example of a metaphor, *argument is war*, as employing a metaphorical frame such that the source domain of war frames the target concept of an argument. According to Lakoff and Johnson (1980: 4), this metaphor is reflected in everyday language by expressions such as:

Your claims are *undefensible*.
He *attacked* every *weak point* in my argument.
His criticisms were *right on target*.
I *demolished* his argument.
I've never *won* an argument with him.
You disagree? Okay, *shoot!*
If you use that *strategy*, he'll *wipe you out*.
He *shot down* all my arguments.

Here, the frame *argument is war* (at least partially) influences how we think and talk about the concept *argument*. This is explained by how attending to certain key elements of the concept war (e.g. *having a position; defending, attacking, retreating; aiming at victory*) makes salient and emphasises certain features of the concept of argument (e.g. *having an opinion; giving reasons in its favour, giving reasons against the opponent's opinion, altering our opinion if the opponent's criticisms cannot be countered; aiming to convince ones' opponent that our opinion is correct*) in ways that help us reason about and with the concept *argument*.

Before proceeding, it is worth acknowledging that it might be hard to hear these last sentences as metaphorical or to paraphrase their communicated content, let alone ascertain that this content is false on a face-value reading. This is to be expected. One study found that widespread metaphors are recognised as metaphors only 17% of the time (Ervas and Ledda 2014). Metaphors can be thought of existing on a continuum, with the strikingly metaphorical sentences uttered by Romeo on one end and the undetected metaphors which feature in Lakoff and Johnson's cases on the other. In between, we have examples of commonly used metaphors which nevertheless are recognisably metaphorical: "Look *how far* we've come", "she's *a pillar* of society", "Danny is *sporting a new bruise*", "the bike was *weaving in and out* of traffic", and many more. These sentences exploit a source domain (physical miles travelled, a pillar holding up a building, wearing an item of clothing, weaving a thread) to shape how we think about a target domain (progress towards a shared goal, playing a positive, pivotal and supportive role in one's community, being proud of a physical bruise, moving between traffic in a certain motion). What all of these metaphor frames have in common, from the highly poetic to the barely detectable, is that they exploit some feature of one domain to illuminate, shape and structure our thinking about another domain.

Although the exact nature of metaphor frames is contentious, theorists of different stripes tend to agree that metaphor frames influence cognizing about a concept by focusing agents' attention to features of the target domain that are paradigmatically instantiated in the source domain. For example, on Camp's (2017) construal, framing involves a restructured characterization of the target domain, one that highlights the features of the target domain that resemble the salient features in the source domain. In a similar vein, Thibodeau, one of the most active researchers on metaphor frames, introduces them as follows: "metaphor frames highlight certain aspects of a target domain and deemphasize others, thereby encouraging specific patterns of inference."

(2017:270). Entman (1993: 52; italics in original) also highlights the way in which metaphor frames call attention to particular features of the target domain: “Framing essentially involves *selection* and *salience*. To frame is to *select* some aspects of a perceived reality and make them more *salient* in a communicating text, in such a way as to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation for the item described.” Finally, Landau et. al. review a wide range of research in social and cognitive psychology that shows how “experimentally manipulating the metaphoric framing of target concepts influences how people perceive and make inferences and judgments about target-relevant information” (2010:1057).

There is ample evidence to suggest that metaphorical frames are both widespread and effective at influencing how we reason about and make inferences involving the relevant concepts. In a series of studies Thibodeau and Boroditsky (2011, 2013, 2015) show that whether crime is framed metaphorically as a beast or a virus influences participants’ evaluations on which crime-prevention measures should be taken, with participants in the ‘crime as a beast’ frame preferring enforcement-based solutions and participants in the ‘crime as a virus’ frame preferring rehabilitation-based solutions. This phenomenon is explained in terms of the metaphor framing how people think about the source domain by emphasising the ‘dangerous’ and ‘in need of a cure’ aspects of crime respectively. A host of further studies have uncovered similar reasoning patterns in other domains. Framing climate change as a war instead of a race correlates with an increased perception of urgency, risk and willingness to change behaviour (Flusberg et al 2017); framing flu in terms of a beast, riot, army or weed as opposed to neutral frames correlates with higher vaccine uptake when compared with neutral frames (Scherer et al 2015); framing talk of the police in terms of guardians as opposed to warriors evokes more positive attitudes towards the police (Thibodeau, Crow and Flusberg 2017); framing ideas as lightbulbs over seeds correlates positively with perceived genius of male inventors, while the reverse is the case for female inventors (Elmore and Luna-Lucero 2017); and personifying changes in stock prices (‘climbing’, ‘slipping’) rather than objectifying them (‘increasing’, ‘decreasing’) correlates positively with people predicting that the current trajectory will continue. These studies suggest that framing a concept in terms of another influences how people think about and reason with it.

Applying metaphor frames to the concept of knowledge, I want to posit a conceptual mapping between the source domain (knowledge on an infallibilist reading) and the target domain (RI+) such that talk and thought about the contents of RI+ are framed at least partially in terms of knowledge. On this view, knowledge talk is metaphorical in the following way: sentences of the form “S knows that P” typically instantiate a metaphorical frame according to which being a reliable informant with regards to P is framed in terms of being infallible with regards to P. This is not to say that speakers typically believe informants to be infallible but rather that aspects of being infallible make salient and bring attention to aspects of being a reliable informant in ways that confer some form of advantage. That is to say, we think about being a reliable informant with regards to P in ways that are at least partially influenced by how we think about being infallible with regards to P.

So, how does appeal to metaphor frames explain how knowledge sentences manage to be useful, despite being infelicitous? It is widely acknowledged that metaphor frames are important heuristics which aid our cognitive reasoning. Specifically, the way in which frames emphasise certain key features while downplaying others helps us to reason about and make

inferences with a given concept, thereby reducing cognitive load (Landau et al 2010, Johnson 2010, Bougher 2012, Thibodeau et al 2017). In Yablo's terminology, metaphor frames often play a presentational role, allowing us to better reason with and about the relevant metaphorically framed concepts. As such, the utility of infelicitous sentences is explained by how framing these sentences in ways that would be infelicitous if taken at face value confers cognitive advantages over presenting the non-face-value content without recourse to a metaphor frame. Put differently, a metaphorically framed content is easier to reason and talk about than the same content expressed literally.

As evidence, there is research to suggest that metaphor frames do not influence reasoning in every instance. They have a markedly reduced effect in cases where the participant has already attained a good level of independent knowledge about the target domain (Reuchamps, Dodeigne, and Perrez, 2018; Robins & Mayer, 2000) and when the metaphor increased rather than reduced ambiguity (Robins and Mayer, 2000). The type of task also effects metaphor uptake, with complex tasks eliciting more metaphor usage than simple tasks such as route learning (Mayer, 1976, Borgman, 1999). This is exactly what one would expect if metaphor were presentationally essential, since simple tasks do not require concepts to be presented heuristically via metaphor. Finally, participants are more likely to employ metaphor to interpret an issue if they are motivated to achieve a certain epistemic status with regards to the issue (Epley, Waytz, and Cacioppo 2007; Landau, Keefer & Rothschild 2014). In a series of studies, Landau, Keefer & Rothschild (2014) showed that participants were more likely to employ metaphors to think about an issue when they had been primed to feel uncertain about the target issue and when they were concerned about it. That motivation towards understanding an issue correlates with metaphor usage is further evidence that metaphor is employed as a tool with which to reason.

A final point worth noting: metaphor frames are essential components of diverse theories on metaphor and feature in all three of the current main perspectives on metaphor (Semino 2008): cognitive (e.g. Lakoff and Johnson 1980), discourse-based (e.g. Cameron et al. 2010) and practice-based (e.g. Reisfield and Wilson 2004). As such, casting the fictional mechanism in the concept of knowledge in terms of metaphor frames does not commit me to any substantive view on metaphor (and indeed, the next chapter explores two different ways of being a metaphor-based epistemic fictionalist)⁵⁰.

Having specified the notion of fiction by which knowledge sentences are directed away from their face-value content and pointed to some general cognitive advantages of metaphor frames, I now turn to the five criteria for a well-motivated fictionalist theory. I argue that epistemic fictionalism so understood meets all five of the criteria. Crucially, it employs the notion of fiction in accounting for how sentences of the form "S knows that P" communicate their non-face-value content, and gives an account of the utility of this redirection. I argue that framing the

⁵⁰ We have seen that Chung's (2020) account employs Lakoff and Johnson's (1980) Conceptual Metaphor Theory (CMT). Metaphor frames are an essential component of CMT. Whereas I have described frames as *influencing* how we think about and reason with concepts, on CMT metaphor frames such as 'argument is war' *structure* concepts, and they do so specifically by using simpler, embodied concepts to structure more complex and abstract concepts. For example, on CMT to be an argument *is to be* an activity structured at least partially in terms of war. CMT therefore applies a stronger notion of metaphor frames than the one I use here. Nevertheless, fans of CMT may read the discussion on how frames influence thought in terms of structuring without too much alteration.

face-value content (RI+) in terms of the face-value content (Infallibilism) confers two distinct advantages: highlighting the demanding standards for knowledge and suppressing irrelevant possibilities of error. These advantages explain why RI+ is typically presented in such a way as to be infelicitous if taken at face value, while remaining useful. That is to say, they explain why knowledge discourse is fictional in the first place.

4.2. Meeting the Fictionalist Criteria

I now argue that a metaphor-frame version of epistemic fictionalism on which the content communicated by sentences of the form “S knows that P” is the non-face-value content RI+ successfully meets the criteria outlined in Chapter 3 for a well-motivated fictionalist theory. Although I briefly discuss each one, I take (i)-(iv) to be relatively straightforward. Most of the space is taken up with defending (v), the requirement that epistemic fictionalism applies the notion of fiction to explain how sentences of the form “S knows that P” succeed in communicating some other, non-face-value content.

i. A theory about sentences in a domain of discourse D

As in Chung’s fictionalism, function-first epistemic fictionalism straightforwardly meets this desiderata. Function-first epistemic fictionalism is a theory about fragments of discourse, specifically, about sentences of the form “S knows that P” and the content that they communicate.

ii. Sentences in D would be infelicitous on a face-value reading

On this view, sentences of the form “S knows that P” would, if taken at face-value, communicate that “S’s evidence for P entails that P”. On the plausible assumption that one’s evidence never, or very rarely, entails the truth of the proposition that it supports, then these sentences are typically false. Furthermore, they are false in ways that speakers can be brought to see with relative ease, by presenting sceptical arguments.

iii. Uttering sentences in D fulfils a useful function

Function-first epistemic fictionalism is especially well-suited to explaining how sentences of the form “S knows that P” fulfil a useful function, for it posits that the function of the discourse is to communicate the contents in RI+. We thus have a ready-made account of the function of re-directing away from the face-value content (“S is infallible with regards to P”) towards a non-face-value content (e.g. “S is a reliable informant with regards to P”). The function is to fulfil a fundamental need – the need to flag reliable informants to one another. Additionally, we already have an account of why the face-value reading cannot fulfil these functions: on a face-value reading, no (or very few) agents count as knowers and therefore none (or very few) of them will be flagged as reliable informants.

iv. Uttering sentences in D fulfils a function in virtue of doing something other than communicating their face-value content

As we saw in Chapter 2, many function-first epistemologists argue against sceptical infallibilism on the grounds that there is no point to an impossibly demanding concept of knowledge. And

indeed, it is surely right that barring widespread error, there would be no point to a concept of knowledge that was exclusively used to communicate its face-value reading, where the face-value reading pegs the standards for knowledge as impossibly demanding. On my proposed view, the function of knowledge talk is to flag reliable informants, points at which inquiry may end, and so on. It therefore serves its function in virtue of communicating something other than its non-face-value content.

v. Applies the notion of fiction in the explanation of (iv)

Unlike Chung's version of fictionalism, function-first epistemic fictionalism also meets criterion 5. First, my view posits metaphor frames as the mechanism by which sentences of the form "S knows that P" come to communicate a non-face-value content. The view therefore centres fiction, and in particular metaphor frames, as the thing which explains how uttering sentences of the form "S knows that P" does something other than communicating "S's evidence for P entails that P". The claim is not that knowledge talk functions in ways that are akin to metaphorical talk, but rather that knowledge talk actually *is* an instance of metaphorical talk. That being a reliable informant is framed in terms of the infallibilist's concept of knowledge is what explains how sentences of the form "S knows that P" serve their function.

Why, then, is this content represented as metaphorical as opposed to being represented at face-value? As was the case for JTB+, the discourse is not plausibly representationally essential: although speakers may not be aware of the content that they use these sentences to communicate, in moments of reflection we are able to paraphrase the non-face-value content into face-value language. I also do not think that the utility of the discourse is best captured in terms of procedurally essential discourse, on which it is indeterminate from the speakers' perspectives whether they wish to communicate the face-value or non-face-value contents. As discussed in relation to Chung's view, it is implausible to suppose that it is indeterminate from speakers' perspectives whether they are infallible with regards to the propositions which they claim to know.

Instead, I want to argue that the knowledge metaphor is presentationally essential. In particular, I want to argue that framing RI+ in terms of infallibilism emphasises the demanding standards for counting as a reliable informant (or ending inquiry, and so on) by highlighting the maximally demanding standards at play in the concept of knowledge. For the function-first epistemologist, attributing knowledge is serious business. The costs of recommending insufficiently reliable informants, of terminating inquiry too soon, of misplaced honour and blame are severe. Terminating inquiry too soon can lead to acting on misinformation at grave personal cost. Inappropriately recommending, honouring and blaming members of the community who are undeserving of recommendation, honour and blame has serious social consequences, for example leading a fellow community member to act on unreliable information or the unjust punishment of a community member. Fulfilling the practical functions of knowledge requires that agents take the high standards seriously. Just as framing the concept *crime* in terms of the concept *virus* makes salient rehabilitation as a crime-management option, so framing *reliable informant* in terms of *knowledge* makes salient the demanding standards necessary for the practice of recommending informants.

One of the key elements of metaphoric frames is how they evoke affective responses. In some cases, metaphors create emotional valences which transfer to the target domain (Bowers and Osborn, 1966; Hitchon, 1997; Johnson & Taylor, 1981; Reali, Soriano, & Rodríguez, 2016; Siltanen, 1981; Thibodeau, Crow, & Flusberg, 2017). Thinking of *flu* as a *riot*, *crime* as a *beast* and *climate change* as a *war* all create powerful affective responses by inducing people to attend to the ways in which the alarming aspects of the source and target domain are similar, and thus feel greater urgency regarding the target domain. There is also some neurological evidence that metaphorical expressions are particularly emotionally evocative, as shown by the increased activation of the emotion-controlling amygdala and anterior portion of the hippocampus in response to metaphorical sentences when compared to non-metaphorical sentences (Citron and Goldberg 2014). Put differently, frames not only influence how we think about concepts, but also influence how we *feel* about them. I think that we see a similar affective result in the knowledge metaphor. Framing *RI+* in terms of *knowledge* elicits a particular feeling, a sense of urgency and anxiety that causes agents to take seriously the tasks of recommending sufficiently reliable informants, of terminating inquiry after sufficient investigation, or honouring and blaming those sufficiently worthy of honour and blame.

Indeed, work on the ever-expanding field of epistemic emotions shows that epistemic emotions such as doubt, anxiety and uncertainty play crucial roles in our cognitive architecture and in mediating our behaviour. Anxiety and doubt motivate us to continue inquiring into an uncertain proposition, while feelings of certainty and the absence of doubt causes us to cease inquiry (Hookway 2002; Nagel 2010; Vazard 2019). Why might this be? Hookway (2002:259) suggests that “doubt of a proposition renders particular questions and lines of inquiry salient”. The emotional reaction of anxiety and doubt renders possibilities of error salient and drives further inquiry aimed at confirming or ruling out these possibilities. The presence of doubt and anxiety, then, highlights the demanding standards for knowledge by making salient certain possibilities of error in need of ruling out. This is precisely what we should expect if knowledge talk were metaphorical with the face-value meaning posited by infallibilism. Metaphor frames generate an emotional reaction which in turn influences action and practical reasoning. In this case, framing reliable informants as infallible creates a heightened emotional reaction (e.g. doubt and anxiety) which motivates action and practical reasoning (e.g. continued inquiry and resistance to attributing knowledge).

One might worry that if this frame really does influence thinking and action, then it is mysterious why we ever attribute knowledge. Why does the knowledge frame not cause us to inflate the standards for a good informant beyond what is needed? In the first instance, the power of frames shouldn't be overestimated. Yes, frames influence how we think about a given concept – but they do not *determine* how we think about a concept. We can, and often do, reject the frame when it has outrun its use. For example, we have seen that framing *crime* as a *beast* (e.g. “crime stalks the city”) leads people to support harsher punishment and less rehabilitation than they would if neutral frames were used. Plausibly this is because the *beast* frame makes salient certain features of *crime*, such as its dangerousness and unpredictability, while suppressing other features such as the possibility of rehabilitation. Despite this pervasive influence of the frame on people's thinking, it is not the case that people would think and talk about *crime* as if it were a *beast* in

literal contexts. They would not, for example, advocate trapping crime by baiting it with honey or recommend pest control technicians take over from the police, even in the presence of metaphor frames.

I've argued that knowledge as a metaphor frame makes salient the demanding standards for knowledge when one is attempting to determine whether a knowledge attribution is warranted. There is still something missing, however. I also want to defend the idea that knowledge as a metaphor frame continues to serve a presentational role even after knowledge has been attributed by making salient that one has done enough, epistemically speaking, to be a reliable informant. In a nutshell, once an informant meets the appropriate standards of reliability with regards to a proposition, framing one's epistemic standing with regards to that proposition in terms of knowledge prevents the needless expending of energy involved in improving one's epistemic position beyond what the situation requires. In these cases the frame makes salient the absence of error typical of knowledge, and uses this to highlight the absence of (relevant) error required of reliable informants. This accounts for why the metaphor frame doesn't typically cause the standards for knowledge to rise above what is needed for recommending reliable informants who fall short of being infallible informants.

As both Craig (1990) and Hannon (2019a) emphasise, fulfilling the functions of knowledge requires not only that knowledge be sufficiently demanding, but also that the standards for knowledge be obtainable. If knowledge were not obtainable, or were obtainable only with great difficulty, then the concept would fall into disuse and would fail to fulfil its function. Function-first epistemologists argue that the practical demands of knowledge push the standards down and away from the infallibilist standards. In this way, function-first epistemologists take themselves to resolve the elusive threshold problem for knowledge. For the function-first epistemologist, the threshold between knowledge and lack thereof consists in the point at which designating a given epistemic position as knowledge promotes the concept's function.

For Hannon (2016, 2019a), the threshold for knowledge and lack thereof tracks the difference between an informant who is sufficiently reliable for the purposes of her epistemic community and one that is not, where a member of the epistemic community are those that it is reasonable to expect would draw on the information. Without a point at which to stop, people would carry on inquiring endlessly and would waste valuable resources seeking increasingly better evidence. However, maximally good evidence is rarely required for practical purposes. We do not reasonably expect the sceptic to draw upon our information, and so the maximally demanding standards imposed by the sceptic do not inform the standards for our practical concept of knowledge. All we need is evidence that is good enough for everyone in our epistemic community – or perhaps evidence that is good enough for our particular projects, where these are more demanding than the communally determined standards.

That knowledge is sensitive to the needs of the community answers the conceptual question of where to peg the standards for knowledge and explains why these standards will not be those posited by the infallibilist. However it leaves unanswered the following psychological question. Even acknowledging that the standards of the community are baked into what it means to count as a knower in that community, we can still ask *why* we stop inquiring once these standards are

reached. Hannon thinks that the standards for knowledge are learned socially. He writes that “socialization and acculturation make us proficient at distinguishing the possibilities that must be eliminated in order to have knowledge from those that typically do not” (2016:617) and that “the practice of knowledge attribution is learned in the same way that humans of sufficient maturity learn to tell when it is appropriate to stop asking “Why?”” (2014:1131) Hannon’s point is not that we stop *because* we take ourselves to be good enough for the community. If it were, then knowledge would not be a motivator on questions that the community do not care about, or where we are investigating something that we never intend to share with anyone else (Kelp, 2011). Rather, Hannon’s claim is that we internalise the standards that the community generally requires given typical concerns and apply them even when not attempting to make ourselves into informants. Hannon (2014) notes that there are good reasons to have one fixed standard. Widely diverging standards would add to cognitive load, as members of the community would struggle to keep track of the various standards in play across various contexts. Further, different standards would require far more context-sensitive information than is practical on the inquirer’s particular project and the informant’s exact epistemic standing before recommending someone as an informant.

I want to suggest that framing the epistemic standard at which one is a reliable informant in terms of infallibilism further diminishes the cognitive load by preventing speakers from dwelling on possibilities of error that are not relevant to the community’s everyday projects. Provided that we are typical members of the epistemic community engaged in typical everyday epistemic projects, sentences of the form “S knows that P” indicate that P is as likely to be right as our purposes require. We need not determine each time whether S is following the shared standards, nor consciously reflect on whether our project is one that the recommender is likely to have considered. Rather, we can assume that S has ruled out all of the possibilities that typical members of the epistemic community (ourselves included) will require. Thinking of being a reliable informant via the frame of knowledge draws attention away from the various possible standards and their relationship to individual inquiries and instead highlights that we can proceed as if there are no relevant possibilities left for us to eliminate. Put in the language of metaphor frames, in sentences of the form “S knows that P”, certain features of the target domain ‘knowledge’ (that there are no $\sim P$ possibilities of error) make salient and focus attention to certain features of being a reliable informant with regards to P (that there are no $\sim P$ possibilities of error that one could reasonably expect to be relevant to members of their epistemic community).

The same point can be made by thinking about the expressive power of non-face-value discourse more generally. Yablo (2000:172) writes that non-literal discourses “show a heartening tendency to boost the language’s power to communicate facts about other, more ordinary, entities”. Similarly, ordinary facts about when an inquiry has gone on long enough and so when one is a sufficiently reliable informant are given an expressive boost by thinking and speaking as if one is infallible with regards to the proposition in question.

To summarise, knowledge is presentationally essential in two distinct, but related, ways. On the one hand, framing what it takes to be a reliable informant via the lenses of infallibilism makes salient the demanding standards for recommending informants, encourages agents to take these

standards seriously and promotes related affective responses. On the other hand, once an informant is deemed sufficiently reliable, the knowledge frame highlights the good epistemic standing that one has in relation to the proposition and discourages further needless inquiry by suppressing considerations about standards and irrelevant error possibilities. The claim that the infallibilism frame captures both of these seemingly competing pressures on the standards for knowledge might strike some readers as suspicious, and may raise the worry that the frame will turn out, in the end, to be practically inert. How are we to implement the frame correctly so that we attend to the relevant feature at the right time, and use it to frame how we think about the target domain?

To begin with, it is important not to overstate this worry. Metaphor is by its very nature highly context-dependent, sometimes picking out one relevant feature of the source domain and sometimes another (Barndern 2007, Camp 2006). As such, interlocutors are often required to pick out the meaning intended by the speaker from among competing possibilities. Consider a traditional example first discussed by Lakoff and Johnson (1980). The metaphor frame ‘*love is a journey*’ is sometimes used to highlight a couple’s shared goals (“Look *how far we’ve come*”, “We’re at a *crossroads*”) and at other times emphasise the difficulties that a couple might encounter in their relationship (“The marriage is *on the rocks*”, “The relationship is *off track*”). Different aspects of the concept *journey* (as involving a shared destination, as involving impediments to travel) frame these different aspects of the concept *love*. Yet, we expect interlocutors to readily distinguish between instantiations of the metaphor frame aimed at making obstacles salient from those aimed at making shared goals salient. Similar considerations apply to the *argument is war* metaphor, which frames both the conditions for winning or losing an argument and the various moves along the way.

Second, the knowledge frame makes different features salient at different times, but these aspects are not *conflicting*. On the knowledge frame, error possibilities correspond to community-indexed relevant possibilities. When one is deciding whether an informant is sufficiently reliable, then the community-indexed relevant possibilities that have not been eliminated are made salient by the knowledge frame, encouraging one to improve their evidence if needed. When one has already determined that the community-indexed relevant possibilities have been eliminated, the frame makes salient that one can rely on the proposition in question and need not inquire further. As is the case for metaphor frames more generally, an internally consistent frame is used to make salient different features of the target concept, and this is possible because different aspects of the source domain map onto corresponding aspects of the target domain. In this case, different aspects of *knowledge* (that it is maximally demanding, that you can rely on it) map onto different aspects of *being a reliable informant* (that it is as demanding as the everyday projects of the community requires, that so long as your project is an everyday project you can rely on an informant’s testimony).

4.3. The Problem of Scepticism

Epistemic fictionalism offers a way of handling sceptical arguments that does justice to both the intuitions that favour scepticism and the intuitions that speak against scepticism. In particular, epistemic fictionalism answers the question: when sceptical worries rear their heads, why do we

concede that we did not actually know some everyday proposition about the external world, yet experience no inclination to stop making knowledge assertions in the future? The answer is that we were never committed to the face-value content of our knowledge claims to begin with. Sentences of the form “S knows that P” communicate the non-face value content that “S is a reliable informant with regards to P”, not the face-value content “S’s evidence that P entails the truth of P”. As such, there is no tension between denying the commitments of a face-value reading of sentences of the form “S knows that P”, while carrying on using knowledge sentences to communicate their non-face-value content.

More carefully, the epistemic fictionalist says that sceptical arguments appear compelling for the simple reason that they are correct when sentences of the form “S knows that P” are given a face-value reading. Consider again the following sceptical arguments:

Sceptical Argument

S1. I don’t know that I’m not a handless brain in a vat

S2. If I don’t know that I’m not a handless brain in a vat, then I don’t know that I have hands

SC. Therefore, I don’t know that I have hands

According to the epistemic fictionalist, the argument is cogent if ‘knows’ is given a face-value reading: for if I know that I have hands, where this means that my evidence entails that I have hands (S1), and I deduce that not being a handless brain in a vat is entailed by my having hands (S2), then I know that I am not a brain in a vat, where my evidence entails that I am not a brain in a vat (SC). Hence, sceptical arguments are compelling because their conclusions are true on a face-value reading.

Crucially, it is a feature of metaphor (and non-literality in general) that interlocutors are always entitled to insist on the face-value reading of uttered sentences, even when doing so would be highly uncooperative (Chung 2018, Camp 2006a). When confronted by an uncooperative interlocutor, speakers are either forced to paraphrase their metaphorical utterance using face-value language, or retract their assertion – on pain of having their linguistic competence regarding the relevant concepts called into question. To use Camp’s (2006a) example, when Romeo tells his cousin Benvolio that “Juliet is the sun”, Romeo can expect Benvolio to understand and respond to the non-face-value content of what he uttered, which might be something like: “Juliet is worthy of worship; that she is the focus of his thoughts and dreams; that she is an exemplar of goodness and beauty; that her goodness and beauty are natural and original, and far exceed those of the other ladies of Verona; and that her goodness nurtures Romeo and helps him to grow emotionally” Camp (2006b:3). But despite this, if Benvolio is feeling uncooperative – perhaps he has grown weary of hearing Romeo complain about yet another love interest – he is entitled to insist upon a face-value reading of Romeo’s utterance. That is to say, he is entitled to reply with something like: “That’s absurd: Juliet isn’t the sun. It’s up there, and she’s a normal girl, here on Earth. It’s older than the Earth, and she’s barely 14 years old. And it’s bigger than the Earth, while her waist is just 26 inches around.”(Camp 2016a:296). Upon being challenged in

this way, Romeo is forced to either rephrase the content that he wished to communicate, and rephrase it using face-value language (e.g. “Ok, fine, I just meant that she’s really beautiful and unique”) or he’s forced to retract altogether (e.g. “Ok, fine, she’s not the sun. But she is really beautiful and unique”). He cannot however dismiss Benvolio’s challenge by insisting upon his own, non-face-value content: he cannot for example reply with “You’re wrong, she *is* the sun”. This is because there is a sense in which Benvolio’s challenge is ultimately correct: on a face-value reading, Juliet is *not* the sun.

Of course, Benvolio’s challenge is highly uncooperative: precisely because Romeo’s utterance is so obviously false, Benvolio cannot seriously have understood Romeo to be attempting to communicate the face-value meaning of his uttered sentence. But even if we judge Benvolio to be a poor conversational partner, we cannot accuse him of linguistic incompetence with the term ‘sun’. It seems rather that Benvolio’s challenge falls well within the remit of competent language use.

As Camp goes on to argue, entitlement to insist on face-value readings becomes even more evident when there are more plausible face-value contents that the speaker could have been taken to mean by a cooperative (as opposed to uncooperative) interlocutor. To use another of Camp’s suggestive examples, suppose Socrates uses the metaphor “Diotima is a midwife” to communicate the non-face-value content that Diotima helps young people work out their philosophical views. Suppose further that Socrates is overheard by a councilman who is investigating, and imprisoning, unregistered midwives. The councilman is perfectly entitled to question Socrates further, demanding to know why Diotima hasn’t registered if she is indeed a midwife. And Socrates cannot simply stick to his intended metaphorical meaning without rephrasing or retracting. Even clarifying that Diotima does not deliver babies is insufficient: the reply “Diotima is a midwife, but she doesn’t deliver babies” in this context does not absolve Socrates of incompetence with the word midwife, unless combined with rephrasing using face-value language or retracting altogether. In contrast, a speaker who uses language at face-value cannot be forced to rephrase or retract by an interlocutor who wishes to switch to non-face-value reading. Someone who speaks at face-value “can accept all the obvious consequences a hearer brings up, or else deny that they follow from what she said, without sacrificing her claim to linguistic competence” (Camp 2006a:298). Camp takes this to show that face-value meanings are lodged into the semantics of metaphorical sentences, while non-face-value meanings are a matter of pragmatics. Regardless of the underlying mechanics, one thing is clear: where there is dispute, face-value content wins out.

That interlocutors can always insist on the face-value reading of an utterance explains why scepticism is so compelling: the infallibilist can always insist on an infallibilist reading of sentences of the form “S knows that P”, forcing speakers to either retract or paraphrase, on pain of incompetence with the word ‘know’ (Chung 2018, 2020). That is to say, although it is perhaps uncooperative to do so, the sceptic can always switch the reading in play to the one on which her sceptical conclusion is true. According to epistemic fictionalism, then, the following two challenges are symmetrical:

Juliet

R: Juliet is the sun

B: That's absurd, she is a girl, not a ball of flaming gas

R: Ok, fine, she's not the sun.

Scepticism

H: I know that I have hands

S: That's absurd, you can't know that you're not a BIV

H: Ok, fine, I don't know that I have hands

Just as Benvolio can insist that both he and Romeo come to read Romeo's initial utterance at face-value, so too the sceptic can insist that both he and his interlocutors come to read an initial claim to know at face-value. In this sense, Fictionalism is a theory that, to use Pritchard's (2018:4739) turn of phrase, "cede[s] the epistemological high ground to the radical sceptic".

Fictionalism not only explains why scepticism is so compelling, but it also explains why the truth of scepticism is ultimately not alarming in everyday life. For in everyday contexts with interlocutors who are not uncooperative, sentences of the form "S knows that P" communicate that S is sufficiently reliable with regards to P. On the assumption that in many contexts it is possible to be a sufficiently reliable informant with regards to P, there will be many contexts in which such sentences are true on a face-value reading. Thus there is a very real sense in which scepticism does not threaten the truth of our knowledge attributions: scepticism doesn't threaten the truth of what we typically communicate when we attribute knowledge. Just as Benvolio's challenge to Juliet not being a flaming ball of gas does nothing to threaten the truth of what Romeo was attempting to communicate - that Juliet is beautiful and unique, perhaps - so the sceptic's challenge that the speaker does not know that she is not a handless brain in a vat does nothing to threaten the truth of what the speaker was attempting to communicate - that they are a reliable informant with regards to their having hands, perhaps. The truth of scepticism, then, presents no reason to cease using knowledge attributions exactly as we do.

Additionally, Chung argues that going fictionalist about knowledge also explains instances of resistance to sceptical arguments. After all, although sceptical arguments are compelling and hard to refute, at least some philosophers have challenged the datum that everyday speakers readily retract their knowledge attributions under sceptical pressures (see e.g. Cohen 2000, Hawthorne 2000, Davis 2007, Fantl and McGrath 2009, Dinges 2016). Rather, these philosophers suggest that everyday speakers are just as entitled to 'stick to their guns' in the face of a sceptical challenge. A good response to scepticism should therefore also account for why some people, at least some of the time, are unmoved by sceptical arguments. Epistemic fictionalism has the resources to do just this: speakers are unmoved because they take themselves to have communicated the non-face-value content of their uttered sentences. That is to say, speakers take themselves to have communicated something that the sceptic has not called into question (that S is a reliable informant with regards to P) and hence that there is no pressure to retract. Indeed, it is widely noted that cooperative interlocutors readily discern, and respond to, what the non-face-value content that speakers intended to communicate, as opposed to the face-value meaning of what

they said (Camp 2006)⁵¹. Speakers going into a conversation expecting a cooperative interlocutor therefore might well experience the sense that they have not misspoken, and that their sceptical interlocutor has failed to engage appropriately. This also explains why sceptical challenges such as brain in vat scenarios can feel a bit silly or like they miss the point: the sceptic is talking past the speaker, addressing the face-value content of their uttered sentence instead of the non-face-value content that they were entitled to expect the sceptic to grasp.

Interestingly, though, even people who are inclined to resist sceptical arguments initially, might well relent and concede that the sceptic gets *something* right. This is consistent with how speakers who use metaphor expecting cooperative interlocutors might update their understanding of the nature of the challenge being presented, especially where the metaphor could not plausibly have been intended to communicate its face-value content.

Compare how A might respond to the following challenge to her metaphorical assertion:

Crime is a Beast

A: Crime is a beast ravaging the city of Addison

B: No, it isn't - crime is an action performed by people

A: Ok, but crime is a beast in the sense that it's savage, dangerous and uncontrollable

B: Why do you keep saying that crime is a beast? Crime isn't a beast, only animals are beasts. Crime is a category for behaviour, not an animal. Animals cannot commit crimes.

A: Ok fine. I didn't mean that crime is a literal beast, I just meant that it's like a beast in being dangerous and hard to control.

Barring some particular reason for why B might struggle to parse metaphorical sentences (let's stipulate that B is not a small child, for example), B can be interpreted as an extremely uncooperative interlocutor. So uncooperative, in fact, that A initially struggles to grasp the nature of B's challenge. However, once A grasps that B is insisting on face-value readings of 'beast', then A grants that B is right - but A grants this on a technicality, while rephrasing and remaining committed to the non-face-value content that they wished to communicate.

I think that a similar pattern of response underlies resistance to sceptical arguments as well. Consider:

⁵¹ This way of presenting how meaning comes apart from 'what is said' is a bit quick. Contextualists about metaphor think that in uttering metaphorical sentences, speakers *assert* the non-face-value content. The point I'm making here, though, does not hang on any particular way of construing how a speaker communicates the non-face-value content of their metaphorical sentences. The point is rather that sceptical resistance originates from speakers expecting their interlocutors to engage with the non-face-value content of their uttered sentence, whether it is part of what is asserted or part of what is pragmatically conveyed.

Bus Stop

A: I know that the bus stop is just around that corner

B: No, you don't - there might not be a bus stop there at all

A: Ok, but know that it's there because I come here often and the bus stop is just next to that pizza shop on the corner

B: Why do you keep saying that you know there's a bus stop there? You can't possibly know that. You might be dreaming, or hallucinating, or you might be in the Matrix where there are no buses and no bus stops.

A: Ok fine. I didn't mean that I know I'm not in the Matrix, I just meant that *provided that we're not in the Matrix*, the bus stop is over there

Just as in the 'Crime is a beast' challenge, the speaker in 'Bus Stop' eventually concedes that her interlocutor gets something right - but what they get right is a semantic technicality, which does not threaten the truth of the communicated content. The sense in which scepticism is correct, then, can be acknowledged without being treated as a genuine source of concern, even in situations where there is initial temptation to stick to one's guns.

I have tried to show how epistemic fictionalism predicts the range of reactions to sceptical scenarios that we find. It explains why sceptical arguments are so compelling (their conclusion is true, on a face-value reading); it explains why even though scepticism is true on a face-value reading, the sceptical conclusion is in no way alarming (scepticism does not threaten the content that we use sentences of the form "S knows that P" to communicate); it explains why some people, some of the time, resist the sceptical conclusion altogether (they expect their interlocutor to discern and respond to what they transparently were attempting to communicate); and finally, it explains the particular pattern of back-and-forth we might expect between the sceptic and a speaker initially inclined to resist (the speaker acknowledges that the sceptic get *something* right, without conceding that they have undermined what the speaker attempted to communicate).

Nevertheless, the function-first epistemologist might protest that going fictionalist about knowledge talk is overkill: for function-first epistemology already has the resources to explain both the resistance to, and appeal of, sceptical arguments. We have already seen how the function-first programme seeks to resist scepticism in everyday life (Chapter 2). However, function-first epistemologists also seek to explain why scepticism gets a handle to begin with, and they do so without appealing to fiction. Craig is explicit in wanting an analysis of knowledge that explains why sceptical challenges are not part of everyday life, while still explaining why they have engaged philosophers in the way that they have. Hannon (more explicitly in 2019a than in 2019b) is also clear about wanting an analysis that does not make sceptical possibilities illegitimate in every possible circumstance.

Craig and Hannon reconcile scepticism with the everyday function of knowledge by confining the legitimacy of sceptical challenges to particular types of inquiry, while still holding on to the

idea that sceptical challenges are illegitimate in everyday contexts. They do this by first acknowledging that certain features of the function of knowledge do push the standards for knowledge upwards towards those posited by the sceptical infallibilist. Craig writes that the process of objectivization conceived as “the tendency to seek to improve the reliability of our information” (1990:117) enables the sceptic to keep pushing how reliable our information must be. The sceptic “extrapolate[s] this tendency without doing anything which the concept of itself forbids; and such extrapolation will lead to scepticism” (1990:117). This process of extrapolation is something that philosophers can legitimately do, and they can do it because the concept of knowledge leaves itself open to being extrapolated in this way; but nevertheless, the motivations of philosophers engaged in this task “cannot possibly be thought to be part of the mental life of any and every human community” (118). Rather, the everyday concept of knowledge serves the needs that do exist in any and every human community, and these practical needs push the standards for knowledge down towards those posited by the fallibilist. Thus Craig distinguishes between the context of the philosopher – “a highly unusual thinker, conducting a highly unusual investigation” (1990:110)- and the everyday context in which knowledge attributions are made, arguing that although we should not expect to find a “hard boundary” between the two, only the former licences the sceptic’s concept of knowledge.

In the same vein, Hannon (2019a) argues that sceptical challenges have a place only in the philosopher’s inquiry conceived of as a “pure inquiry” (Williams, 1978), in which the philosopher undertakes the radical project of seeking truth and avoiding error without antecedently assuming the truth of any of her beliefs. This is Descartes’ quest – to start entirely afresh, and inquire with no regard for the practical constraints that typically govern our conduct in inquiry. In contrast, inquiry in everyday contexts strives for “balancing our need for truth (and the avoidance of error) with the constraints and limitations that we fallible humans encounter in daily life” (Hannon 2019a:218). For Hannon, then, the challenges raised by the sceptic are appropriate uses of the concept of knowledge, but only in the context of the philosopher’s particular project of pure inquiry. This successfully “explains why the skeptic’s argument is plausible and yet not a threat to ordinary knowledge” (Hannon 2019a:220), and does so without appeal to fiction. So, the function-first epistemologist might protest that a function-first realist analysis of knowledge already possesses the features that I have claimed as a distinct advantage of a function-first epistemic fictionalist view.

In response, I agree with Craig and Hannon that the sceptic exploits a legitimate part of the concept of knowledge – but disagree that the legitimacy of sceptical challenges is confined to particular, and unusual, philosophical mode of inquiry. On my view, the sceptic exploits the face-value reading of sentences of the form “S knows that P”, while everyday inquiries trade on their non-face-value meaning.

I think my account has two advantages over Craig and Hannon’s line of argument. First, on Craig and Hannon’s view it is difficult to see why the sceptic would ever succeed in propelling us into a philosophical inquiry, when we are already engaged in a practical inquiry. Yet, sceptical challenges demand a response, even when they are raised in everyday contexts. I have argued that when confronted with a sceptical scenario, the options are: rephrase, retract, or resist. Even granting that at least sometimes it’s appropriate to stick to one’s guns, we are still compelled to

acknowledge that the sceptic gets something right. Why would this be, if the sceptic's challenge is only appropriate in the context of philosophical inquiry, a project far removed from the present practical inquiry? The epistemic fictionalist can answer this question. The sceptic exploits the face-value meaning of sentences of the form "S knows that P", and so is always entitled to insist on her preferred reading. The speaker must engage with the face-value reading whenever her interlocutor demands it. Despite this, engaging with the face-value reading does not alter the nature of the practical inquiry at hand, which is why the speaker can simply rephrase and carry on with her practical inquiry. Epistemic fictionalism, then, accounts for why sceptical challenges are, at least in one sense, perceived as needing to be engaged with even in everyday non-philosophical contexts.

The second problem for Craig and Hannon's view is that it is very difficult – perhaps, impossible – to locate a principled distinction between philosophical and practical contexts, a point which both Craig and Hannon acknowledge. For his part, Craig is sympathetic to the idea that philosophical inquiries are also practical. He asks whether the "practically minded man" would apply the methods characteristic of pure inquiry, in which every assumption is questioned, and concludes that "[q]uite likely he would, unless prescient as well as practical" (1990:118). Taken from the perspective of the inquirer starting out on their inquiry, it is not apparent whether one's inquiry is going to be practical or pure; for the answer will in part depend on how successful the inquiry turns out. Descartes, famously, took his inquiry to have succeeded in yielding foundations for the sciences that are "stable and likely to last" (1641:6) – what goal could be more practical than that? Furthermore, if we define pure inquiry as inquiry that challenges all of our assumptions, then philosophical inquiry *doesn't* qualify as pure inquiry. Philosophers too inquire from within a programme, holding steadfast some assumptions in the course of their inquiry⁵². But epistemic fictionalism doesn't need to demarcate types of inquiries in this way. For the fictionalist, sceptical challenges target the face-value reading of sentences of the form "S knows that P", whether they appear as part of practical or philosophically-minded inquiries.

Nevertheless, this does raise an interesting possibility, which would require further investigation to make compelling. One possible reason for sceptical challenges being perceived as more legitimate in the philosophy seminar than in everyday contexts is that one of the key projects in philosophy centres on uncovering the face-value meaning of everyday sentences. Many philosophers take the search for truth to be one of the principal aims of philosophy⁵³. And often what philosophers ask are questions like: "what is beauty?", "what is morality?", "what is free will?". In practical terms, philosophers address these questions by considering whether sentences of the form " ϕ is beautiful", " ϕ is moral", " ϕ was a freely performed action" are true. Furthermore, philosophers tend to be interested in whether such sentences are *really* or *fundamentally* true, beneath appearances. One natural interpretation of what it takes for a

⁵² And indeed, sceptical challenges are perhaps even more frustrating in non-epistemic philosophical contexts than they are in everyday contexts, precisely because philosophers engaged in investigating ethical or aesthetic considerations plausibly take themselves to be assuming that those possibilities do not obtain.

⁵³ Brennan (2010), Dietrich (2011), Horwich (2012) and Chalmers (2015). Although cf: Brandom (2001) Beebe (2018), and Hannon and Nguyen (2023)

sentence to be “really” true is that it be true on a face-value reading⁵⁴. If these tentative suggestions are along the right lines, then the fictionalist too can capture the sense in which sceptical challenges are more *relevant* to philosophical contexts: philosophers are often using, or investigating, the face-value meaning of fragments of discourse. Or at least, they are doing so more often than is typical of everyday contexts. Crucially, though, the fictionalist holds that although sceptical challenges are more *relevant* in philosophical contexts (on account of speaking to the very thing under investigation) they are *legitimate* in all contexts (on account of speaking to the face-value reading of sentences of the form “S knows that P”).

4.4. Advantages of Fictionalism

In this section, I want to show how the fictionalist view that I have proposed yields several benefits. Specifically, I wish to show how the arguments that speak in favour of sceptical infallibilism that I canvassed in chapter 1 work for fictionalism too.

4.4.1. Epistemic Possibility and CKAs

Epistemic fictionalism provides a neat explanation of why CKAs sound infelicitous: to come out true, they would need to mix face-value and non-face-value meanings. But metaphorical sentences cannot be mixed in this way – if an interlocutor wants to insist on a face-value reading, then they successfully force the sentence as a whole to take on a face-value meaning. As discussed, the sentence “Diotima is a midwife, but she does not deliver babies” is infelicitous in a context where an interlocutor is insisting on a face-value reading of the second clause (i.e. the interlocutor has made it clear that they want to find out whether Diotima delivers literal babies), precisely because insisting on a face-value reading of second first clause carries over to the first clause. Furthermore, specifying that Diotima doesn’t deliver babies gives interlocutors to understand that they are employing a face-value reading, since whether or not someone delivers babies for a living is a function of whether or not they are a (literal) midwife.

In the same way, a CKA such as “I know that Alana committed the murder, but she might not have” is infelicitous if an interlocutor is insisting on a face-value reading of “she might not have”, where the ‘might’ is a function of whether the speaker infallibly knows that Alana committed the murder. Furthermore, in mentioning the possibility that Alana might not have committed the murder, the speaker gives the interlocutor to understand that they are employing a face-value reading. Thus, the first clause is false on this reading.

Additionally, if the sentence as a whole were taken metaphorically, then it would also come out false: to give “I know that Alana committed the murder” a non-face-value reading is to say that I am a reliable informant with regards to the question of whether Alana committed the murder. But the metaphorical mapping on which sentences of the form “S knows that P” communicate that S is a reliable informant with regards to P, maps onto related notions such as possibility. What makes “S is a reliable informant” true is that S has eliminated all of the possibilities of error

⁵⁴ Indeed, many fictionalists describe their core commitment that a sentence is not true *on a face-value reading* as the claim that the sentence is *not true* (see for example Liggins 2012, Eklund 2019, Caddick Bourne 2018).

that are salient or relevant to the community. Hence, what makes a non-face-value reading of sentences of the form “S knows that P” true is whether S has eliminated these possibilities. Extending the metaphor to the second clause, then, would be to treat the clause “she might not have” as communicating that there are some salient or relevant possibilities that S has not eliminated; thus making the sentence as a whole false on a non-face-value reading.

4.4.2. Knowledge does not admit of degrees

In a similar vein, on epistemic fictionalism, knowledge does not admit of degrees. There is only one degree of knowledge when sentences of the form “S knows that P” are read at face-value, and it is that posited by the infallibilist. On a non-face-value reading, there are multiple contents that the metaphor frame “S knows that P” might be used to communicate. For example, it might communicate that S is a reliable informant with regards to P, or that S has ended inquiry into whether P, and so on, depending on what the sentence is being used to communicate.

Furthermore, the non-face-value reading is fundamentally context-sensitive, in the sense that which strength of epistemic position must be obtained in order for S to count as a reliable informant with regards to the question of whether P will depend on the situation. Chapter 6 delves into more detail regarding the context-sensitivity of metaphorical knowledge. For now, note that following Hannon (2013), we can expect that situations where the stakes are particularly high will call for a better epistemic position than situation where nothing much hangs on the truth of P. Features of the context, then, are part of what makes the non-face-value reading of sentences of the form “S knows that P” true. Nevertheless, this is not to say that there are different degrees of knowledge. For in each instance, what makes the non-face-value reading true or false is whether S is sufficiently well positioned to be a reliable informant.

4.4.3. Knowledge is Qualitatively Different to Non-knowledge

Epistemic fictionalism is well placed to explain the difference in quality between states that may be described with metaphorical sentences of the form “S knows that P” and those which may not be. The simple answer is that one may metaphorically attribute knowledge only when something of practical utility has been achieved. Say that following Craig (1990) and Hannon (2019a), knowledge picks out the point at which a potential informant is a sufficiently reliable informant with regards to P. Insofar as this distinction is one that it is practically useful to make, then there is a qualitative difference between knowledge and non-knowledge (see Hannon 2014). Similar considerations apply to the other functions of knowledge: insofar as stopping inquiry is useful, then there is a qualitative difference between a point which may metaphorically be described with knowledge talk and a point which may not be, and so on for other posited functions of knowledge talk.

However, going fictionalist about knowledge also answers an objection to function-first epistemology as a general method. Function-first epistemology has been criticised for giving an inadequate account for why the concept of knowledge has emerged in the first place (Feldman 1997, Rysiew 2012). What does the concept add, over and above the concept of a reliable informant? Without an answer to the question, the function-first epistemologist struggles to

explain the emergence of the concept of knowledge by appeal to social and practical needs. As Rysiew (2012:278) puts the point, “we have at our disposal other epistemic terms we can and do use in picking out good informants—‘trustworthy’, ‘reliable’, ‘always right’, etc. But given that there are plenty of terms available for picking out informants of one or another degree of reliability what’s special about ‘know(s)’?” Answering this question is necessary for answering the question of what distinguishes knowledge from non-knowledge.

Craig’s answer is that knowledge is objectivized, while reliable informant is relative to specific needs and projects. Someone who is unwilling to disclose their information, for example, is a knower but not a reliable informant. Building on this line of argument, Hannon’s answer is that ‘knows’ marks that particular degree of reliability such that anyone in the community may draw on her information for their everyday projects (2019:45). Therefore, there is something special about knowledge that these other epistemic terms do not capture. Still, we might wonder why, on a fundamentally social function-first view, these other epistemic terms are not likewise indexed to the needs of the community. A potential informant who is unwilling to disclose their information may still be reliable, where to be reliable means to meet the epistemic standards needed for the community’s projects, despite the fact that inquirers lack access to their reliable information. Why not think that the terms ‘reliable’ and ‘trustworthy’ already pack in the notion of being sufficiently reliable and trustworthy for the community’s projects? To echo Rysiew, what’s special about knows?

The fictionalist has a ready-made reply: knows is special because it is the only epistemic term that frames being sufficiently reliable in terms of infallibilism. Thus, it is the only epistemic standard that confers the cognitive benefits discussed in the previous section: helping agents to take seriously the demanding standards for knowledge and shielding them from excessive doubt once these standards have been reached. As such, the fictionalist response can bolster the function-first claim that no other epistemic term is capable of fulfilling the distinct role occupied by knowledge. Further, fictionalism can do so without giving up the central function-first claim that the social and practical functions of knowledge are the primary explanation for why knowledge fulfils these roles.

4.4.4. Knowledge is Closed under known entailment

Finally, epistemic fictionalism affords a principled explanation for the competing intuitions surrounding the principle that knowledge is closed under known entailment. That is to say, fictionalism is able to diagnose both steadfast acceptance of, and rejection of, the principle of closure. Acceptance of the principle of closure comes from considering the face-value reading of sentences of the form “S knows that P”. For on a face-value reading, the principle of closure holds in ways that are uncomplicated: the face-value reading of sentences of the form “S knows that P” is that “S’s evidence for P entails that P”. It follows that if P entails Q, then S’s evidence for P also entails that Q. That when given a face-value reading, knowledge is closed under known entailment thus explains the intuition that closure is a cornerstone datum that theories of knowledge need to account for. As noted, this principle is what lurks behind the sceptical arguments like the one discussed in §4.3, and repeated here for convenience:

Sceptical Argument

S1. I don't know that I'm not a handless brain in a vat

S2. If I don't know that I'm not a handless brain in a vat, then I don't know that I have hands

SC. Therefore, I don't know that I have hands

On the other hand, the principle that knowledge is closed under known entailment looks much less plausible when focusing on the non-face-value reading of sentences of the form "S knows that P". Taking the 'recommending informants' function, it seems as if S could well qualify as a good informant with regards to the question of whether P, while failing to qualify as a good informant with regards to the question of whether Q, even though P entails Q and S believes that Q on the basis that P entails that Q. In 'Sceptical Argument', if "know" is taken to indicate being a reliable informant, then S1 might well be true on a non-face-value reading, while SC. might well be false⁵⁵. Although she posits different non-face-value readings to the ones that I have proposed, Chung (2020) thinks that the epistemic fictionalist ought to embrace closure failure on non-face-value readings, and indeed that this is the very thing which explains why sceptical arguments fail to threaten everyday knowledge attributions. That knowledge might not be closed under known entailment on non-face-value readings, then, explains why many philosophers have ultimately been satisfied with rejecting the principle of closure.

Furthermore, the function-first epistemic fictionalist is particularly well placed to account for why in many cases, one can extend one's knowledge from P to entailed premise Q through competent deduction, even though knowledge is not strictly speaking closed under known entailment on a non-face-value reading. And the reason is that for most everyday propositions Q, where P entails that Q and where S has competently deduced Q from P, then S being a reliable informant with regards to P also makes S a reliable informant with regards to Q. Recall that function-first epistemology analyses concepts by reflecting on how that concept might serve its purposes in typical cases in which the need that the concept has emerged to fulfil is present. Thus, counter-examples which are out of the ordinary do not speak against an analysis. In most ordinary situations, then, everyday users of knowledge talk can fruitfully extend the pool of propositions which they may legitimately describe themselves as knowing via competent deduction from P to entailed proposition Q.

⁵⁵ Perhaps whether closure holds on non-face-value readings depends on what precisely the non-face-value reading is. If sentences of the form "S knows that P" communicate the non-face-value content that S may stop inquiry into whether P, then perhaps both S1 and SC are true, since SC is a paradigmatic example of an inquiry that it would be fruitless to continue. This is particularly salient in cases of multiple premise closure: it is difficult to see how one could stop inquiry into whether P, stop inquiry into whether Q, and stop inquiry into whether R, and yet continue to inquire into whether $P \wedge Q \wedge R$. For surely to inquire into whether $P \wedge Q \wedge R$ just is to inquire into whether P, into whether Q and into whether R. Thus it seems that if sentences of the form "S knows that P" do communicate the content that "S may stop inquiry into whether P" then knowledge perhaps is closed under known entailment, even on a non-face-value reading.

Function-first epistemic fictionalism, then, is well placed to explain why intuitions on whether knowledge is closed under known entailment pull in different directions. In philosophical contexts, the principle of closure (like the sceptical arguments that they underpin) appears intuitive when sentences of the form “S knows that P” are given a face-value reading, and counter-intuitive when they are given a non-face-value reading. In everyday contexts, that principles of closure captures a frequently reliable way of competently extending the pool of propositions that one may legitimately describe themselves as knowing. This explains why closure at least sometimes appears to hold in everyday contexts, while failing to hold in others.

4.5. Conclusion

In this chapter, I have proposed and fleshed out a new theory of knowledge: function-first epistemic fictionalism. According to function-first epistemic fictionalism, sentences of the form “S knows that P” admit of (at least) two readings: a face-value reading, which expresses the content predicted by the sceptical infallibilist; and a non-face-value reading, which communicates contents related to the functions that knowledge plays, for example, that S is a reliable informant with regards to P. This is the content that sentences of the form “S knows that P” are typically used to communicate. In particular, sentences of the form “S knows that P” are underpinned by a metaphor frame: the source domain of knowledge understood along infallibilist lines frames the target domain of what it takes to be a reliable informant with regards to a proposition (or to end inquiry, and so on). I have argued that this way of going fictionalist about knowledge talk meets the criteria for a well-motivated fictionalist theory outlined in §3.1.4, in particular, that it accounts for how elements of fiction are crucial to knowledge talks serving its functions.

The main advantage of epistemic fictionalism is how it handles sceptical arguments: sceptical arguments are sound on a face-value reading of sentences of the form “S knows that P”, which explains why sceptical arguments are compelling; but they nevertheless do not threaten the truth of what such sentences are typically used to communicate, which explains resistance to sceptical arguments. Thus epistemic fictionalism is a view on which sceptical infallibilism is (at least in one sense) correct, but on which the truth of sceptical infallibilism is not a problem for everyday speakers of knowledge talk. I have also argued that epistemic fictionalism has the resources to explain the intuitions surrounding concessive knowledge attributions; the sense in which knowledge is qualitatively different from non-knowledge; and the intuitions on the principle of closure. The next two chapters aim to defend the view that I have proposed by defending it from objections (chapter 5) and by comparing it to rival views (Chapter 6).

Chapter 5: Objections and Replies

This chapter responds to objections to epistemic fictionalism. The first half (§5.1) responds to what is often considered the most pressing objection to fictionalism in general: that it is phenomenologically implausible. Specifically, that fictionalism is at odds with what everyday users of the target discourse take to be the content of their utterances; their psychological attitude towards the content of their utterances; and the type of speech act that they engage in when making these utterances. I highlight the striking parallels between the debate between competing versions of fictionalism and the debate in metaphor theory about how metaphor succeeds in communicating its non-face-value content, with a view to showing how different metaphor theories can ground their counterpart version of fictionalism. I argue that versions of both content and force fictionalism have the resources to respond to the phenomenological objection. In doing so, one of my aims is to show how there is more than one way of constructing a plausible epistemic fictionalism, and that both content and force versions of the view are available. I also address the issue of dead metaphors (§5.2), and respond to the objection that fictionalism cannot account for the origins of the target discourse's face-value content (§5.3).

5.1. The Phenomenological objection: knowledge talk doesn't *feel*/like fictional talk

Many critics of fictionalism take it as a strike against fictionalism that it contradicts speakers' phenomenological experience of using the target discourse (See Stanley 2001; Eklund 2007; Chrisman 2008, Reynolds 2009, Armour-Garb and Woodbridge, Brock; Friend 2008, Liggins 2010). After all, in most domains where fictionalism is posited, it does not seem to speakers and hearers that they are engaged in non-face-value talk: users of moral discourse do not feel like they are speaking in non-face-value ways when they say "it is morally impermissible to torture", and users of knowledge discourse do not feel like they are communicating in non-face-value ways when they say "I know that P".

Here are some representative formulations of this phenomenological objection:

Competent users of arithmetical discourse will certainly deny that they are pretending when they discuss arithmetic. In such cases, the hermeneutic fictionalist must maintain that the fact that the language user is pretending is not accessible to her, even in principle. (Stanley 2001:46)

When we say that the number of starving people is very large, we do not feel ourselves to be talking (even a bit!) about the content of a mathematical story. Our subject matter is *people* and our thought is that a lot of them are starving. (Yablo 2001: 76)

If I think that you are completely out of line, I am not making believe that you are completely out of line, even as I thump the table, go purple in the face, march you out of the door and forbid you to show your face again. (Blackburn 2005:33)

Mathematical discourse just doesn't *seem* at all like Cowboys and Indians (Eklund 2007)

It just doesn't seem to me that I am operating under some pretence when I make a moral claim, and, more generally, I doubt that ordinary speakers have the phenomenology of pretence when it comes to moral discourse. (Chrisman 2008:3)

We do not seem to be pretending when we engage in the discourses of which pretence theories are claimed to be true: indeed we seem not to be pretending. For instance, if we introspect, it seems to us that we believe, rather than pretend, that $2 + 2 = 4$. (Liggins 2010: 768)

In short, it does not seem to speakers and hearers that they are engaging in non-face-value talk when they utter sentences in the domains of discourse which the fictionalist takes typically to be uttered in a non-face-value spirit. Insofar as contradicting speakers' phenomenological experience counts against a theory, then the phenomenological objection is a serious objection to fictionalism in all domains, including to epistemic fictionalism.

As discussed, there are two broad views which fall under the banner of 'fictionalism': force fictionalism and content fictionalism⁵⁶. According to force fictionalists, utterances of sentences in the target discourse are not assertions, but constitute some other speech act, for example, the speech act of pretending to assert (Joyce 2020:117), of quasi-assertion (Kalderon 2005:112), or of launching a game of make-believe (Walton 1993:53). In pretending to assert that P, the speaker succeeds in communicating or drawing attention to some other content, or perhaps some non-propositional communication⁵⁷. On content fictionalism, a typical utterance of a sentence in the target discourse constitutes an assertion. However, the assertion does not communicate its face-value content. Rather, it communicates a different, non-face-value content⁵⁸. Although they need not be, these views can be combined. For example, a fictionalist could hold that in uttering a sentence in the target discourse domain one pretends to assert the face-value content, and thereby succeeds in asserting the non-face-value content. The metaphor fictionalist theories advanced by Yablo and Joyce are of the force variety. In particular, they both hold that in uttering sentences in the target discourse domain, speakers pretend to assert the face-value content of the sentence. In contrast, Chung's metaphor fictionalism is of the content variety. On Chung's view, in asserting "S knows that P" one succeeds in asserting one (or more) of multiple candidate non-face-value contents.

Both versions of fictionalism are susceptible to a potential phenomenological objection. The fictionalist says that utterances of sentence S in discourse D are typically non-face-value utterances, either in that they are non-assertoric speech acts (force fictionalism) or in that they assert a non-face-value content (content fictionalism). Yet, it seems that speakers do not take themselves to be making non-face-value utterances, but rather take themselves to be asserting the

⁵⁶ These roughly map onto what in chapter 3 I called assertoric and non-assertoric fictionalism. For the sake of consistency with the literature, this chapter uses the content / force terminology.

⁵⁷ See Joyce (2005) and Kalderon (2005) on moral talk, Yablo (2000) on mathematical talk, Crimmins (1998) on attitude ascriptions, Kroon (2001) on identity talk, Armour-Garb and Woodbridge (2019) on proposition and truth talk.

⁵⁸ See Lewis (1978) on meta-fictive discourse, Rosen (1990) on possible world talk, Field (1989) and Balaguer (1998) on mathematical talk, Brock (2002) on fictional character talk.

face-value content of their utterances. The fictionalist, then, posits that speakers are simply mistaken about whether they are engaged in non-face-value or face-value talk. According to the phenomenological objection, the error that fictionalism saddles everyday speakers with is implausible: we do expect, and should expect, that speakers not make the types of mistakes that fictionalism posits. Indeed, in typical cases of engagement in non-face-value discourse, it seems that speakers are aware that they are engaged in non-face-value discourse. For example, speakers acting on a stage or reading a bedtime story plausibly consciously pretend to assert their lines, and speakers engaged in non-face-value talk such as exaggeration or sarcasm are aware that their utterances communicate a content other than the face-value content. This does not appear to be the case for speakers engaged in the discourse domains that the fictionalist wants to treat as non-face-value, such as moral talk, modal talk, mathematical talk and knowledge talk. That fictionalism seems to posit implausible error regarding whether a speaker is engaged in non-face-value talk counts against fictionalism.

In fact, there are three related aspects of how the fictionalist analyses the target discourse which appear to contradict speaker phenomenology. Fictionalism makes predictions that contradict speaker phenomenology as regards the attitude that speakers bear towards the content of their utterance of a sentence in the target discourse; as regards the type of speech act constituted by an utterance of a sentence in the target discourse; or as regards the content of an utterance of a sentence in the target discourse. Not all versions of fictionalism generate phenomenological objections on all three counts, but both force and content fictionalism are each susceptible to at least one version. I briefly discuss each objection in turn, before turning to the metaphor fictionalist response.

First, force fictionalism contradicts speaker phenomenology as concerns the attitude that speakers bear towards the content of their utterances. The fictionalist says that the speaker bears a particular attitude towards the content of their assertion (e.g. quasi-belief, acceptance, pretence), despite the fact that speakers cannot determine via introspection that they possess this attitude and worse, despite the fact that it seems to speakers that they bear an entirely different attitude towards the proposition, typically, belief. Stanley argues that this is a problem for fictionalism because it is implausible to suppose the existence of a psychological attitude, such that one can bear this attitude towards a proposition and yet it be entirely inaccessible to them that they do so. Here is Stanley (2001:47):

If the hermeneutic fictionalist is correct, then *x* can bear the propositional attitude of pretense toward a proposition, without it being in principle accessible to *x* that *x* bears the propositional attitude of pretense toward that proposition. But this introduces a novel and quite drastic form of failure of first-person authority over one's own mental states.

Not only is it inaccessible to *x* whether they have this attitude towards the proposition communicated by sentences in *D*, but worse: speakers typically take themselves to bear the attitude of belief towards these propositions. As Brock (2014) makes explicit, there are two sides to the objection from the phenomenology of attitude argument: on the one hand, there is something lacking in everyday phenomenology, to wit, the sense that we are engaging in fiction

of some type; on the other, there is something present in everyday phenomenology, to wit, the sense that we believe the content of our utterances. The second error seems to involve agents in a more radical and implausible form of introspective failure. As Brock (2014:582) puts it, “it is one thing to suggest that an agent may be in a mental state of which she is unaware; it is quite another to suggest that she is aware (or seems to be aware) of being in a mental state that she is in fact not in. While higher order theories of consciousness leave room for the bare possibility of an appearance/reality distinction of this kind, one does well to resist positing such illusions willy-nilly”.

The fictionalist might question the extent to which such failure in attitude is “novel and drastic”, as Stanley claims. After all, there are domains where agents’ attitudes are opaque to them. For example, unconscious bias is a fairly uncontroversial example of speakers possessing attitudes that they are unaware of possessing. However, unlike cases of bias, speakers who utter sentences in the discourse domains given fictionalist treatments tend to behave in line with the attitudes that they explicitly express. Speakers who utter sentences of the form “it is morally impermissible to ϕ ” and “there are three cats under the table” typically also behave as if they bear the attitude of belief towards the proposition “it is morally impermissible to ϕ ” and “there are three cats under the table”. That fictionalism contradicts how speakers feel about their own attitudes, then, does seem to be an unusual result that is not explained away by positing a mismatch between what speakers take their attitudes to be and what their behaviour suggests their attitudes are.

Relatedly, force fictionalism appears to contradict how speakers feel about the type of speech act that they are engaging in when they utter sentences in the target discourse. The force fictionalist posits that we do not assert the content of uttered sentences, but rather engage in a non-assertoric speech act that only resembles assertion in virtue of its surface structure. Yet, it seems to speakers that their utterances of sentences in the target discourse domain constitute assertions. Speakers who utter sentences of the form “it is morally impermissible to ϕ ” take themselves to be asserting that it is morally impermissible to ϕ , and speakers who utter sentences of the form “S knows that P” take themselves to be asserting that S knows that P. Terrance Cuneo (2006) argues that many widespread theories about what constitutes a given type of speech act hold that speaker intention is a key condition for producing a given type of speech act. Specifically, Cuneo argues that producing a non-assertoric speech act requires that the speaker endeavours to produce a non-assertoric speech act³⁹. One might follow Grice (1957) in claiming that to produce a given speech act, the speaker must intend to produce some effect in the hearer (e.g. that the hearer come to believe a given proposition), and for this effect to be the result of the hearer recognising the speaker’s intention to produce this effect. For assertoric speech acts, the speaker might intend the hearer to come to believe the content of the speaker’s assertion, and to do so by means of recognizing the speaker’s intention for the hearer to form the belief. For non-assertoric speech acts, the speaker might intend the speaker to come to recognise a command, request or compliment. Alternatively, one might follow Searle (1969) in claiming that different speech acts have distinct purposes and communicate distinct psychological states. In the case of assertion,

³⁹ Cuneo aims this objection at expressivism, which also posits that utterances of sentences in the target discourse are non-assertoric speech acts.

the speaker communicates the psychological state of belief towards the content of the assertion, and in doing so commits themselves to the world being a certain way, and opens themselves up to blame should it turn out that it is not the case. In contrast, the point of directives is to get hearers to act in certain ways, and communicates the psychological state of desire⁶⁰.

On either of these pictures, what the speaker intends directly bears on what speech act she produces. Part of what it is to engage in a non-assertoric speech act is to wittingly do so: a speaker who intends to produce belief in her hearer, or intends to take responsibility for the world being a certain way, is making an assertion. A speaker with other intentions is not making an assertion. Yet force fictionalism holds that speakers regularly succeed in making non-assertoric speech acts despite not intending to perform a non-assertoric speech act, and despite actively intending to perform an assertion. Therefore, force fictionalism appears to be incompatible with widely accepted theories on the constitutional conditions for performing a non-assertoric speech act.

The content fictionalist can avoid these particular objections. According to content fictionalism, speakers do have the attitude towards the content of their utterances that they take themselves to have, namely, belief; and they do produce the speech act that they take themselves to be producing, namely, assertion. So, content fictionalism does not contradict speaker phenomenology as regards the speech act produced and as regards attitude towards the content of their utterance. However, the content fictionalist runs up against the third phenomenological objection. According to content fictionalism, speakers use sentences in the target discourse to communicate something other than its face-value content. Yet it appears to speakers that they are using the discourse to communicate a face-value content. The content fictionalist, then, contradicts speaker feelings on what the face-value content of their utterance is, and whether they are communicating the face-value content of their utterances. As Yablo points out in the above quote, speakers engaged in number talk (e.g. “the number of starving people is very large”) do not feel as if the contents of a mathematics story is the subject of their uttered sentence. Similarly, speakers engaged in moral talk or knowledge talk also feel that their utterances concern features of the world, and not the contents of a story.

5.1.2. Metaphor and Fictionalism

Unlike other fiction mechanisms employed by fictionalists, metaphor offers a prima facie promising avenue for resisting all three versions of the phenomenological objection. After all, unlike the phenomenology of playing Cowboys and Indians (Stanley 2001:38), telling a story (Brock 2014:577) or watching a play (Joyce 2020:112), the phenomenology of engaging in metaphorical talk does not differ (or at least, does not drastically and noticeably differ) from the phenomenology of using literal, face-value talk. Metaphor fictionalism, then, has a head start when it comes to answering the phenomenological objection compared to fictionalist theories which posit that the target discourse is an instance of, or is analogous to, these other, phenomenologically distinct forms of engaging in fiction. According to the metaphor fictionalist,

⁶⁰ Note that this is the case even if the speaker is being insincere. The speaker who utters “I promise to pay you back tomorrow” intends the hearer to take the utterance as a reason to think that the speaker intends to pay the hearer back tomorrow. This is the case even if the speaker has no intention to repay the hearer.

speakers who utter sentences such as “it is morally impermissible to ϕ ” no more intend to commit themselves to the face-value content that there are attitude-independent moral facts such that it is morally impermissible to ϕ , than speakers who utter sentences such as “Juliet is the sun” intend to commit themselves to the face-value content that Juliet is a blazing ball of plasma. Rather, the same process (or a relevantly similar process) of re-direction that enables Romeo’s interlocutors to retrieve his intended meaning – something along the lines of ‘Juliet is beautiful and worthy of worship’ – enables interlocutors engaged in moral talk to retrieve each other’s intended meanings.

The nature of this process of re-direction is a matter of ongoing debate in metaphor theory. Interestingly, this debate closely mirrors the debate between the content and force fictionalist, with metaphor theorists disagreeing about whether metaphorical sentences are assertoric or non-assertoric speech acts. On some metaphor theories, speakers who utter metaphorical sentences do not assert the non-face-value content of their metaphorical utterance, but perform a different speech act altogether which communicates the non-face-value content. This speech act might be the act of pretending to assert (Hill 1997, Walton 1993), or implicating (Grice 1957, Camp 2006). On some views (e.g. Davidson 1978, Searle 1979) a metaphorical utterance asserts the face-value content, with the purposes of conveying the non-face-value content. The metaphorical utterance might ‘call to mind’ (Searle 1979:104) or cause us to ‘notice’ (Davidson 1978:37) salient similarities between the metaphor’s source and target domains. Call these theories ‘force’ metaphor theories. Like force fictionalist theories, force metaphor theories posit that the non-face-value content of the target sentence (in this case, metaphorical sentences) is not asserted. In contrast, what we might call ‘content’ metaphor theories hold that speakers who utter metaphorical sentences directly assert a content that differs from the face-value content. Most contemporary content metaphor theories take this asserted content to be the non-face-value content (e.g. Stern 1985, Bezuidenhout 2001, Wearing 2006), so that uttering “Juliet is the sun” is to assert that “Juliet is radiant”. But this is not the only option: Aristotle, for example, took the content asserted by metaphorical utterances to be identical to the content communicated by the corresponding simile, so that Romeo’s utterance asserts that “Juliet is like the sun”⁶¹.

Since both content fictionalism and force fictionalism find analogues in metaphor theory, both versions of fictionalism are in principle amenable to at least some versions of metaphor fictionalism. Whether a particular metaphor fictionalist theory falls under force or content fictionalism will depend on which metaphor theory underpins it⁶². This choice also determines

⁶¹ Aristotle writes that the simile “is a metaphor, differing from it only in the way it is put; and just because it is longer it is less attractive.” [The Rhetoric, 1410b] Roberts (Trans)

⁶² In my view, the choice of metaphor theory on which to build one’s fictionalist theory should be driven by independently plausible views on metaphor. Armour Garb and Woodbridge (2015) argue that philosophical fictionalism need not be beholden to existing positions in the philosophy of fiction. That is to say, moral fictionalism need not be modelled on the fiction mechanism that the theorist believes underpins literary fiction, but need only posit that moral discourse employs some fiction mechanism. This mechanism might be different to the mechanism underpinning literary fiction. I agree with Armour Garb and Woodbridge on this point, but the case for the metaphor fictionalist in particular is different. The metaphor fictionalist wants to use *metaphor* as the mechanism to explain the redirection that takes place in the target discourse. She does therefore need to be beholden to what she takes to be the best theory of metaphor in general. Otherwise, we lose the sense in which her fictionalism is *metaphor*

which version of the phenomenological objection the metaphor fictionalist theory will be required to address. I will argue that both force metaphor fictionalism and content metaphor fictionalism offer promising responses to their respective phenomenological objections. However, uncovering how metaphor fictionalism might address the phenomenological objection requires a serious look at how metaphors work, which is in itself a debate with no clear consensus. Although some metaphor theorists (e.g. Wearing 2006, Walton 1993, Hills 2017) do remark upon the consequences of their views for philosophical fictionalist theories, fictionalists do not typically consider how their view might be altered in light of findings from metaphor theory⁶³. Yet, the underlying metaphor theory does matter a great deal for metaphor fictionalism: if one's analysis of the target discourse domain is that utterances of sentences in that domain are typically metaphorical, then what it means for a sentence to be 'metaphorical' determines what it means to be a metaphor fictionalist about that domain.

The remainder of this 'phenomenological objection' section strives to achieve two aims: (1) show how different metaphor theory underpinnings yield different versions of metaphor fictionalism, and (2) argue that some (but not all) of them are well suited to respond to the phenomenological objection. In particular, I argue that pretence theories of metaphor are not good candidates for underpinning a phenomenologically unproblematic fictionalist theory. Throughout the discussion, I move between discussing objections to the metaphor theory underpinning each version of metaphor fictionalism, and the resulting fictionalist view itself. For epistemic fictionalism to work, it's important that both the underlying metaphor theory and the resulting fictionalist theory also work: if the underlying metaphor theory fails *as a general metaphor theory* then we lose the sense in which the resulting fictionalism is metaphor fictionalism, and disconnect knowledge from metaphorical talk. And if the metaphor theory works on its own terms but yields an unworkable epistemic fictionalism, then this should be taken as evidence that knowledge talk is not metaphorical talk after all (or alternatively, as evidence that despite appearances, the metaphor theory is not correct after all).

Since several force fictionalists base their fictionalism on pretence (e.g. Yablo 1998, 2010, Toon 2016, Joyce 2018) I now turn to discussing pretence force metaphor fictionalism in detail (§5.1.3), before moving on to non-pretence force metaphor views (§5.1.4.) and finally, content metaphor theories (§5.1.5)

5.1.3. Pretence Force Metaphor Fictionalism

Joyce (2018) and Yablo (1998) defend a version of pretence metaphor fictionalism that is modelled on Kendall Walton's (1993) influential prop-oriented pretence metaphor theory. On Walton's view, metaphor is just one example of a wider class of speech acts: prop-oriented

fictionalism. This is not a mere classification cost. Positing that the target discourse is underpinned by a mechanism other than metaphor returns the 'metaphor' fictionalist to having to explain why their target discourse is underpinned by a *sui generis* mechanism; and barring some added explanation, blocks them from claiming that the phenomenology of engaging in the target discourse is the same as the phenomenology of metaphor. Furthermore, such a fictionalist could not use the arguments from chapter 6 which rely on metaphor frames.

⁶³ A notable exception is Chung's (2020) knowledge metaphor fictionalism, based on Lakoff and Johnson's (1980) Conceptual Metaphor Theory (CMT).

pretence, in which speakers engage in pretence as a means of communicating something about the real world. In Walton's taxonomy, prop-oriented pretence contrasts with content-oriented pretence, in which speakers engage in pretence for the sake of engaging in the pretence itself. To use a classic example, two children riding their bikes might play a game in which they pretend that their bikes are horses. They might say: "my horse is winning the race!" or "Your horse is going too fast - wait for me!" as part of their game. The children's interest is in the pretence itself - they care about the game, and engage in content-oriented pretence for the purposes of playing the game. Now imagine a parent watching the children play, and worrying that they really are cycling too fast given their proximity to a busy road. When the parent says "Your horse is going too fast - slow down!" they are, in a sense, engaging in the game - at least to the extent that they are engaging in the rule according to which bicycles are horses. But the parent is not doing so for the purposes of playing along. Rather, the parent wants to communicate something about the real world, and specifically, about the bicycle: that it really is going too fast. The parent thus engages in prop-oriented pretence aimed at communicating facts about the real world. Uttering sentences about horses in this context typically succeeds to communicate content about bicycles via what Walton calls *principles of generation*. These are the conditions that specify what the real world must be like in order for the face-value content to be true in the game. For example, in the example at hand, the principles of generation state that a bicycle is a horse, so that "your horse is galloping faster than mine" is true in the game if your bicycle is going faster than mine, and false otherwise.

On Joyce's (2018) reading of Walton, metaphor involves this type of prop-oriented pretence. A speaker pretends to assert the face-value content of their metaphorical utterance (e.g. Romeo pretends to assert that "Juliet is the sun") and engages in this pretence in order to communicate something about the real world (e.g. that Juliet is radiant). Crucially, it need not be phenomenologically accessible to speakers that engaging in metaphor means engaging in prop-oriented pretence. Indeed, Joyce argues that the very fact that there is an ongoing academic debate about whether metaphor involves pretence suffices to demonstrate that engaging in pretence is not always marked by a distinct phenomenology. For if it were evident by introspection whether uttering a metaphor involves engaging in pretence, then there would be no academic debate - we could introspect and resolve the question by noting the results.

Joyce argues that reflecting on how metaphor involves pretence without users experiencing the phenomenology of pretence affords a moral fictionalist drawing on a pretence theory of metaphor the ammunition with which to address the phenomenological objection. Joyce is careful to stop short of claiming that moral discourse *is* metaphorical: rather, "the claim is that reflection on how metaphor works helps dispel certain challenges" (2020:109). Put differently, his point is that whether a domain of discourse involves pretence is a "theory laden truth" (Joyce 2020:1154) to be determined by discourse analysis, not introspection. Whether metaphorical or moral discourse involve pretence is a question to be settled by carefully considering metaphorical and moral discourses, not by adjudicating on the accompanying phenomenology of using them. If this is right, then the absence of a 'pretence' phenomenology when engaging in a discourse domain does not cut either for or against pretence theories that posit pretence in that discourse

domain. The aim of Joyce's argument is therefore to undercut the phenomenological objection altogether.

I disagree with Joyce about whether metaphor not being marked by a pretence phenomenology helps the fictionalist who does not wish to claim that their target discourse domain is metaphorical. The fictionalist would need to give some reason to think that her target discourse is relevantly similar to metaphor, or at least some independent reason to think that the type of fiction involved in moral discourse is also one that lacks an accompanying pretence phenomenology. However, even setting aside this disagreement, there are reasons to be unpersuaded by pretence metaphor theories and consequently, reasons to be unpersuaded by pretence fictionalist theories that use metaphor as the model for responding to the phenomenological objection. In particular, Camp (2006) and Wearing (2006) have mounted a series of objections to the claim that uttering metaphorical sentences is to unwittingly engage in a speech act of pretence.

To start, Camp objects that pretence theories of metaphor mischaracterise metaphor frames. According to pretence theories of metaphor, in uttering "Juliet is the sun", Romeo "pretends that she just plain is exactly that" (Hills 1997:147), and if the metaphor is apt, this act of pretending that Juliet actually is the sun calls attention to the real circumstances that make it possible and worth engaging with the pretence that Juliet is the sun. But metaphors do not function by getting speakers and hearers to wholesale imagine that the target domain 'just plain exactly is' the source domain. Pretending that Juliet is the sun would be to imaginatively imbue her with as many characteristics of the sun as possible, while stripping her of any of her real characteristics which she does not share with the sun. But this is not what metaphor frames do - rather, a metaphor frame encourages selective attention to particular features of the target domain *as the target domain*. Romeo's hearers are invited to continue to think of Juliet as Juliet, but to use the sun metaphor to highlight, introduce or restructure how they conceptualise key aspects of Juliet. Hearers who tried to wholesale imagine Juliet as possessing as many of the sun's attributes as possible would completely misinterpret Romeo's intended meaning. Pretence, then, would limit rather than enhance metaphor's ability to fulfil its function of using a source domain to frame a target domain. And as discussed, if pretence is not the right account of metaphor, then a pretence fictionalism built on pretence metaphor underpinnings will be of no use.

Worse, given that pretence metaphor theory was supposed to help resolve the phenomenological objection to fictionalism, Camp argues that the phenomenology of pretence is incompatible with the phenomenology of metaphor. When speaking metaphorically, subjects can (although, they need not) be aware of the target and the source domains as distinct, separate objects. When engaging in pretence, attending to the target domain as it in fact is gets in the way of the act of pretending, so that focus must remain exclusively on the source domain. Hearers trying to pretend that Juliet is the sun will be unable to also focus on Juliet as she actually is, without interrupting the pretence. Yet when we use metaphors, our attention typically is on the target domain: on Juliet herself. This would be bad news for the pretence fictionalist, who would have avoided the phenomenological objection at the level of fictionalism only to re-encounter it at the next level of analysis, at the level of the metaphor theory underpinning the fictionalist one.

In response, Walton (2013) accuses Camp of badly misconstruing his proposal, arguing that properly understood his own proposal resembles Camp's allegedly competing one much more closely than it resembles Camp's reconstruction of his view. In particular, Walton complains that his proposal aims to capture precisely what Camp claims that Walton misses: the way in which metaphors exploit the source domain to communicate about the target domain. Rather than asking people to imagine that the target domain is the source domain, Walton argues that on his view, engaging with the source domain generates rules which establish what the target domain must be like if it is to be true, in the fiction, that the target domain is the source domain. Here is Walton (2013:4):

The make-believe games metaphors introduce are not ones in which certain fictional truths are given, established as it were by stipulation. They are games consisting in conditional principles of generation, which determine what is or would be fictional should the prop possess certain properties. If "Your computer has a virus" is fictional in a speaker's implied game, this is so not because the speaker asks hearers to imagine it, but because, and only if, the computer has certain (difficult to specify) properties which, given the conditional principles, generate this fictional truth. The speaker's claim is that the computer does possess these properties. If the speaker is right, only if she is right, it is fictional in the game that the computer has a virus.

So according to Walton, no attitude of pretence is needed to successfully engage with pretence fictionalism. Hearers interpret metaphorical sentences by recognising the game of make believe invoked by the metaphorical utterance, and consequently come to understand what would need to be the case for the target domain to (fictionally) be the source domain. Crucially, 'what would need to be the case' is highly context sensitive. On the understanding that viruses (literally) can only attack living things such as humans or animals, and further understanding that computers are not living beings, that your computer is a living being is not part of what needs to be the case for "your computer has a virus" to be true. What does make the sentence true according to the rules of the implied fictional game is that your computer contains a piece of code or programme which aims to alter how a computer operates and which is designed to spread from one computer to another. When the speaker utters "your computer has a virus" it is this second meaning that she intends to convey.

As far as Walton's own theory is concerned, this response, if accepted, would dismantle the disagreement between himself and Camp. However, this does little to help the pretence (force) metaphor fictionalist. For once we drop the pretence aspect of Walton's metaphor theory, it is unclear precisely what type of speech act he takes metaphorical utterances to be. Sometimes Walton writes as if utterances of everyday metaphors are assertions, which would make his view a content fictionalist view. For example, he writes that part of what metaphors do is expand our repertoire of concepts to include concepts of being 'metaphorically ϕ ' (e.g. metaphorically laughing, metaphorically back to square one, metaphorically the sun) and that "in uttering a metaphor one may assert that some such concept applies in a certain instance" (53), which suggests that the non-face-value content is at least sometimes asserted. Later, Walton also writes that in cases where the concept is unclear, it is likewise unclear what (if anything) the speaker asserts, which also suggests that when the concept is clear the speaker asserts it. However, in the

same paper, Walton makes comments which could be taken to support a non-pretence force fictionalist view. He writes that utterance of a metaphorical sentence “implies or suggests or introduces or calls to mind a (possible) game of make-believe” (46); that it might “invoke, remind us of, the game of make-believe” (48); that it “may be an act of verbal participation” (46), and in the following paragraph that “some metaphorical utterances may not be assertions at all” (46). This is not intended as a criticism of Walton’s account. A real virtue of Walton’s view is its sensitivity to the different types of metaphor, from literary to mundane, and the sense in which they are a nonhomogeneous category. I think that the most charitable interpretation of Walton’s remarks about assertion is that uttering mundane metaphors asserts the non-face-value content, in virtue of the fact that the metaphorical concepts generated by the relevant game are obvious, while uttering novel, unfamiliar, literary or otherwise vivid metaphors might be to perform a non-assertoric speech act such as implying or introducing a set of principles of generation. Nevertheless, the point remains that interpreting Walton’s view as either (or perhaps, as *both*) a view on which metaphorical utterances are assertions, or as a view on which they are non-assertoric implicature means that his view is not one that invokes the attitude of pretence, and therefore is not a view which posits that agents have an attitude of pretence towards the content of their utterances. Therefore, Walton’s metaphor theory is no good to the pretence fictionalist.

Leaving aside whether Walton's view properly counts as a version of pretence force fictionalism, there are other reasons to think that pretence metaphor theories are ultimately misguided, and so ill-suited for underpinning fictionalism. Wearing argues, convincingly in my view, that positing pretence does little to explain how hearers come to interpret a metaphorical utterance. She argues that pretence views do not make sense of the process by which pretending helps to uncover the non-face-value content. Specifically, hearers need to work out which aspects of pretending are informative for retrieving the speaker's meaning. For example, precisely which aspects of pretending that Romeo has asserted that Juliet is the sun assists in retrieving the content Romeo intends to convey? To answer this question, the hearer would first need to know what rules of generation govern the 'x is the sun' metaphor. So, if the hearer is not already familiar with the metaphor, she needs to work out the principles of generation in order to work out how pretending is supposed to reveal the non-face-value content. But working out the principles of generation just is to work out the metaphor’s non-face-value content. If the hearer can work out the principles of generation that she would need to work out in order to usefully pretend, then she can work out the non-face-value content directly. In other words, positing pretence seems superfluous to explaining how hearers retrieve content.

This problem is made particularly salient in the case of twice-apt metaphors, which are true at face-value, in addition to being true metaphorically. Examples include sentences such as “there is a storm coming” uttered when there are dark clouds on the horizon in addition to some non-weather-related foreseen future events, and “No man is an island”, which is (trivially) true at face-value but also intends to communicate the non-face-value content that no person is disconnected from other people. Camp and Wearing both independently argue that twice-apt metaphors pose a particularly acute version of the discussed problem for pretence views: since these sentences are true at face-value, it is hard to see what hearers are supposed to pretend. It cannot be that hearers pretend that the face-value content is true, either in reality or in the game, since one

cannot pretend what one knows to be the case. As Camp (2006:110) puts it, “because these sentences already are true, there would be no imaginative work left to do”. At the very least, any imaginative work will not take the form of pretending that the face-value content is true in the game.

In response, Joyce argues that Wearing and Camp misidentify the relevant act of pretence. On pretence theories of metaphor, one pretends *to assert that* no man is an island, one does not pretend *that* no man is an island. Just as an actor playing Hamlet can pretend to assert “It is very cold” even if it is in fact cold in the theatre, so a speaker using a twice apt metaphor may pretend to assert the metaphorical sentence, even if the face-value content is true. Still, this reply does not explain what it is about pretending that is supposed to help hearers understand the speaker’s meaning. Pretending that the speaker has asserted the non-face-value content tells us nothing about what to pretend - as the case of twice-apt metaphors show, it is not by imagining what conditions would need to obtain in order to make the sentence true at face-value. Therefore, pretence metaphor theory is, in my view, not equipped to underpin a force fictionalist view capable of withstanding the phenomenological objection to fictionalism.

5.1.4. Non-Pretence Force Metaphor Fictionalism

A better solution for the metaphor force fictionalist, in my view, is to abandon the pretence element of force metaphor fictionalism altogether and instead argue that metaphorical sentences belong to a different and more familiar class of speech acts: conversational implicature. On this view, to say that “Juliet is the sun” is not to assert, but to communicate via Gricean implicature that Juliet is beautiful, radiant, etc. This move would enable the force fictionalist to build their fictionalism on a metaphor theory that does not involve pretence (thus avoiding the problem of explaining how one can pretend without experiencing the phenomenology of pretence), but which nevertheless classes metaphorical talk as non-assertoric. On this view, to claim that sentences of the form “S knows that P” are metaphorical would be to claim that such sentences are not asserted, but that they nevertheless do not involve the phenomenology of pretence.

Grice’s own view of metaphor subsumes metaphor under implicature. Grice argues that metaphorical expressions flout the maxim of quality. Hearers recognise that speakers have said something false, and that speakers are trying to be co-operative and therefore not intending to say something manifestly false. Therefore, hearers go looking for alternative meanings. Here is Grice (1975:53) on metaphor:

Examples like you are the cream in my coffee characteristically involve categorical falsity, so the contradictory of what the speaker has made as if to say will, strictly speaking, be a truism; so it cannot be THAT that such a speaker is trying to get across. The most likely supposition is that the speaker is attributing to his audience some feature or features in respect of which the audience resembles (more or less fancifully) the mention substance.

Many metaphor theorists, including Camp, have followed Grice in treating metaphorical utterances as communicating their non-face-value content via implicature. Camp (2006:280)

sums up the Gricean implicature view as holding that "speakers who speak metaphorically say one thing in order to mean another". The Gricean implicature view of metaphor enjoys several advantages as a theory of metaphor: it subsumes metaphor within a larger and widely accepted speech act of implicature; it gives a plausible explanation of the relationship between face-value and non-face-value content; and unlike the pretence view, it seems to get the order of interpretation the right way round: speakers first register the face-value content of the sentence, then look for an alternative meaning, and directly retrieve the non-face-value meaning from the sentence, without employing an intermediate pretence step.

A fictionalism built on these foundations seems promising for responding to the phenomenological objection which afflicts force fictionalism, for on this view metaphor belongs to the widespread category of implicature. Metaphorical speech (and hence a fictionalism built on it) therefore does not involve 'fiction' in the standard sense. It does not involve pretence or make-believe, for example. Rather, we should expect the phenomenology of metaphor to resemble the phenomenology of engaging with implicature more generally. Consider the following typical example of implicature:

A: Are you going to the department party on Friday afternoon?

B: I have to prepare next week's lectures.

B has implicated, but not asserted, that she will not come to the party. But we should not expect B's utterance to be accompanied by specific phenomenology, and certainly not by the same phenomenology as when engaged in acts of fiction like storytelling or playing a game of make-believe. Likewise, on the Gricean Implicature view of metaphor, we need not expect a particular phenomenology to accompany metaphor use.

Despite this, a critic of a fictionalism underpinned by an implicature metaphor theory might object that even conceding that this view explains the absence of a 'fiction feeling' when engaging in the target discourse, it does not speak to the problem of fictionalism disrespecting speakers feeling as if they are engaging in the speech act of direct assertion. For in other examples of maxim flouting, speakers can readily be made to see how their meaning was not asserted. For example, returning to the above example, A might press B to be specific by asking "Does that mean that you're not coming?". Metaphor is likewise highly amenable to such clarifications. Benvolio can press Romeo by asking "Do you mean that Juliet is radiant?". Indeed, in typical cases of metaphor, speakers can readily be made to "feel the distance" (Yablo 2000:222) between the semantic meaning encoded in the uttered metaphorical sentence and the meaning that they were attempting to convey. In contrast, it is unclear whether speakers engaged in the discourse domains analysed by fictionalism do feel that there is any distance between what they assert and what they mean to communicate. Nevertheless, one can perhaps argue that many of the domains of discourse analysed by fictionalists are ones where such distance is felt - for example, Chung (2018) argues that sceptical arguments and their success in engendering knowledge retractions is evidence of a felt distance between the meaning encoded in sentences of the form "S knows that P" and the meaning that they typically express.

Another objection to this type of fictionalism targets the underlying implicature metaphor theory. The objection is that implicature metaphor theory makes the wrong prediction regarding how speakers and hearers respond to metaphorical sentences. In particular, Bezuidenhout (2001) argues that interlocutors directly agree or disagree with the non-face-value content of the utterance of a metaphorical sentence as if had been asserted⁶⁴. Bezuidenhout (2001:156) illustrates this point with the following dialogue, in which three colleagues are discussing who should be the next chair of their department:

A: How about Bill?

B: Bill's a bulldozer.

A: That's true. But isn't that a good thing in this case? We want someone who'll stand up to the administration and get things for our department.

C: I disagree that he's a bulldozer; that exterior hides someone who's basically insecure. But, either way, Bill wouldn't make a good chair

In this interaction, the truth of the face-value content of B's utterance is not under discussion: everybody agrees that Bill is not a literal bulldozer. The disagreement revolves around the non-face-value content of B's utterance. Bezuidenhout (2001:157) writes that "the most natural construal to put on such a dialogue is that B says something, that B says it with assertoric force, and that A and C are either agreeing or disagreeing with what B says". In contrast, the metaphor implicature view holds that A and C are agreeing or disagreeing with what B has merely implicated, as opposed to asserted. But this doesn't seem right: in straightforward cases of implicature, one cannot signal agreement or disagreement with the face-value content in order to signal agreement or disagreement with the non-face-value content. The student in the classic letter writer case cannot defend herself from the implication that she is bad at philosophy by replying "I disagree, my handwriting is terrible!". This is because typically one cannot disagree with what a speaker implies (that the student is bad at philosophy) by disagreeing with what they assert (that the student has good handwriting).

In response, Camp objects to the view that implied content is never part of what interlocutors respond directly to, and attributes this widespread misconception to the tendency in philosophy to treat as paradigmatic instances of implicature where speakers are being deliberately less than obvious in what they intend to communicate. Instead, Camp argues that it is appropriate to treat speakers as having asserted "contents to which they have *openly and obviously* committed themselves by their utterance" (2006a:286, original italics). Therefore, whether it is appropriate to treat an implicature as asserted or not depends on how obvious the speaker intends to be in communicating the implied content. A speaker who says, "I could use some salt on this pork chop" can be reported as having asserted that they would like some salt on the pork chop and a speaker who sarcastically says "What a brilliant idea: let's spend our last dollar on beer! Then I suppose we can hitchhike home" can be reported as having asserted that we should save money

⁶⁴ See also Hill (1997) and Stern (2002)

for the taxi ride home⁶⁵. Yet, Camp claims that in both of these examples, the gap between what the speaker means and the sentence's conventionally encoded content is so great that the communicated content cannot plausibly be considered part of what is asserted. Rather, they are examples of implicature – and specifically, an implicature so obvious that it can be reported as, and responded to, as if the implied content had been directly asserted. Camp argues that the non-face-value content of metaphor likewise “falls on both sides of the line between what it is and is not appropriate to report someone as saying” (2006:286). In the case of commonly used metaphors, it is plausible that speakers are intending to openly and obviously commit themselves to having asserted the non-face-value content of their utterances and so it will be appropriate to directly respond to the implicated content. Contrastingly, highly poetic metaphors are notoriously unobvious and difficult to paraphrase, and so it will typically be inappropriate to respond to a speaker as having asserted a particular non-face-value content through utterances of these metaphors.

Metaphor fictionalism posits that sentences in the target discourse domain are typically communicated via standard metaphors, and would therefore count as an ‘obvious’ metaphor in the sense that the communicated content is obvious⁶⁶. We should expect, then, for metaphorical sentences in moral and knowledge discourse to be appropriately reported as asserted, and directly responded to as if they were asserted. It follows that we can perhaps even expect them to be accompanied by the phenomenology typical of assertion, even though communicating and retrieving this communicated content occurs via the mechanism of implicature⁶⁷. Treating metaphor as an instance of implicature, then, can ground a force fictionalist theory with the resources to respond to the phenomenological objection. For on this view, the communicated metaphorical content is treated as asserted, and it is therefore unsurprising that speakers feel as if they have engaged in assertoric speech acts when they utter sentences in the target discourse.

5.1.5. Content Metaphor Fictionalism

As discussed, the content fictionalist faces a slightly different phenomenological objection. The content fictionalist posits that sentences in the target discourse communicate a different content than the one encoded in their face-value reading. Specifically, the content fictionalist posits that some underlying fictional mechanism converts the content asserted by utterances of sentences in the target discourse, so that a non-face-value content is asserted. Menzies and Pettit (1994) write that sentences uttered in the course of telling a story do not commit speakers to the truth of those sentences because the utterance is understood as containing an elliptical expression of the

⁶⁵ Both examples from Camp (2006:285)

⁶⁶ Crucially, this does not mean that it be obvious that a sentence is metaphorical.

⁶⁷ One might worry that the implicature view of metaphor makes knowledge fictionalism superfluous. On this view, the metaphor theorist claims that implicature is the mechanism underpinning how saying that “S knows that P” communicates that “S is in a good enough epistemic position with respect to P”. But we might wonder why appeal to metaphor, for it seems that we could generate this result entirely with face-value talk combined with traditional Gricean implicature. The letter writer who praises the student’s handwriting as a means to imply that they are a poor philosopher uses entirely face-value talk. In the same way, a sceptical infallibilist might hold that “S knows that P” is uttered at face-value, but used to implicate that S is in a sufficiently strong epistemic position (e.g. Dodd 2011, Russell 2021, Bonjour 2010). I return to this objection in chapter 6.

sentence “In the story, ...”. Fictionalists have argued that various discourse domains are likewise best interpreted as involving a tacit operator or prefix. For example, Rosen’s (1990) modal fictionalism posits that claims about possible worlds are elliptical for claims about the content of a story, namely, the story according to which modal realism is true⁶⁸. Brock’s (2002) analysis of fictional character discourse posits that critical sentences (e.g. “Holmes symbolizes mankind’s ceaseless striving for truth”) are implicitly prefaced by an “according to realist’s hypothesis” (Brock 2002:9) prefix. Field (1989) and Balaguer (2009) argue that story operators are what determine the truth of sentences in the target discourse domains, so that the truth of mathematical sentences is determined with reference to ‘story of mathematics’, which Balaguer (2009:131) spells out as “the claim that there actually exist abstract mathematical objects of the kinds that platonists have in mind”⁶⁹. Balaguer also analyses propositional talk in the same way, so that whether a proposition is true is determined by whether it is ‘true-in-the-story-of-propositions’ (1998:809). Story operators aren’t the only version of content fictionalism. Indeed, Chung’s own version of epistemic fictionalism is a content fictionalist view. Although she leaves open which mechanism underwrites the content redirection, on Chung’s view sentences of the form “S knows that P” communicate a non-face-value propositional content, albeit it is not always determinate which content is being communicated.

The content fictionalist respects speaker phenomenology with regards to the type of speech-act engaged in with an utterance in the target discourse, and with regards to the attitude that speakers bear towards the propositions communicated by their utterance; but it transgresses against speaker phenomenology by positing error regarding what their sentences are about. For example, speakers engaged in modal talk (e.g. “Paul might be at the party tonight”, “there might have been blue swans”) take themselves to be saying something about the real world – about what, in this world, is possible or probable, and not about a story in which modal realism is true. Neither does it help to restructure the sentence so that the speaker’s intended subject is the subject of the sentence (e.g. “Paul is such that, according to possible-world realism, there is a world compatible with our evidence in which Paul comes to the party”). For a start, many speakers are unfamiliar with possible world theories of modality, and do not take themselves to be referencing this theory at all. Secondly, as Armour-Garb and Woodbridge (2015) point out, the reconstruction misses what is important about what the speaker is trying to communicate. In the example at hand, the speaker is trying to communicate about what their evidence suggests with regards to Paul’s coming to the party - this aspect of their communication is missed in the paraphrase. Furthermore, some domains of discourse do not even admit of such paraphrases. For example, mathematical fictionalism posits that mathematical entities such as numbers are fictive entities, and so there is no non-fictive entity to act as the subject of a reconstructed sentence. For example, on mathematical fictionalism, no true sentence can take the form of “the number eight is such that,

⁶⁸ On Rosen’s version of modal fictionalism, modal discourse concerns what is the case in the story according to which David Lewis’ modal realism is true. Rosen concedes that other theories of possible worlds could ground competing fictionalist theories, and suggests Armstrong’s (1986) as a candidate alternative.

⁶⁹ Balaguer’s view is somewhat different to the other versions of content fictionalism, because he holds that mathematical sentences communicate their face-value content. What makes the view fictionalist is that the truth conditions of these sentences are dependent on the mathematical story, and not on the communicated content.

according to the story of mathematics, it is double four”; because ‘the number eight’ falls outside of what the ‘according to the story of mathematics’ operator quantifies over. Since mathematical fictionalism considers all numbers fictive, the fictionalist about mathematical talk cannot hope to restructure sentences that involve number entities to preserve the sense in which the numbers are themselves subjects of these sentences.

There are views in metaphor theory that align well with the spirit of content fictionalism (and hence that make ideal underpinnings for a content metaphor fictionalism) in that they posit that metaphorical sentences are typically assertoric. These metaphor theories likewise posit that uttering metaphorical sentences such as “Juliet is the sun” directly asserts the non-face-value content “Juliet is radiant” (Stern 1985, 2000, Bezuidenhout 2001, Wearing 2006, Hills 1997)⁷⁰. On these views, the non-face-value content is encoded in the metaphorical sentence as one of its multiple meanings.

Josef Stern’s view in particular closely resembles the ‘story operator’ versions of fictionalism. Modelled on Kaplan’s analysis of demonstratives containing a ‘Dthat’ operator, Stern argues that the content communicated by metaphorical sentences implicitly contains a silent ‘Mthat’ operator as part of its (grammatical) logical form. Stern’s formulation says that “For every context *c* and for every expression φ , an occurrence of ‘Mthat[φ]’ in a sentence *S* in *c* (directly) expresses a set of properties *P* presupposed to be metaphorically associated [m-associated] with φ in *c*”⁷¹. For example, the logical form of “Juliet is the sun” is “Juliet Mthat [‘is the sun’]. Mthat[‘is the sun’] contributes the set of properties which are metaphorically associated with ‘is the sun’ in the context of utterance (e.g. being radiant, beautiful and nourishing). In doing so, it contributes towards the predicate picking out the proposition ‘Juliet is radiant, worthy of worship, the centre of Romeo’s world’ as the content of the sentence. Crucially, for Stern the operator is part of the grammatical structure of the sentence, so that the non-face-value metaphorical content is directly asserted. Precisely which content is picked out by the predicate to which the Mthat operator applies is highly-context sensitive, and a wide range of factors might contribute to a property becoming m-associated with a source domain. Further, some m-associated properties do not have face-value paraphrases. Finally, it’s worth noting that Stern’s metaphor theory operates against the backdrop of a thoroughly context-sensitive theory of linguistic meaning in general. For Stern, every sentence (including face-value sentences) admit of multiple meanings, and hearers disambiguate between possible contents each time they encounter an utterance.

Operator metaphor views like Stern’s are closely related to story operator content fictionalist views, and hence naturally suggest a way of spelling out the mechanics underpinning a content fictionalist view. Crucially, a fictionalism based on an operator metaphor theory affords a promising way of answering the phenomenological objection. For on operator metaphor views like Stern’s, it is not apparent to speakers and hearers (at least, not without careful reflection) that the logical form of the sentences that they assert features an operator. Rather, speakers communicate, and hearers pick up on, the content communicated by the set of properties

⁷⁰ Hills’ view combines force and content fictionalism, in that it posits pretence but also that the non-face-value content is asserted.

⁷¹ This formulation is Wearing’s (2006:15) abbreviated version of Stern’s (2000:115) formulation.

presupposed to be metaphorically associated with the expression indicated by the Mthat operator. I want to suggest that understanding story operators as akin to Mthat operators, which determine which non-face-value content is asserted by an utterance of sentences that contain it in their logical form, explains how story operators contribute to non-face-value content production. To see how this might work, consider one of Yablo's (2000:214) examples of metaphor:

1. She's got the smarts

On Yablo's analysis, smarts-talk is relevantly similar to maths-talk in that face-value readings of smarts-talk seem to commit speakers to an ontology on which 'smarts' entities exist, that agents either possess or do not possess, while a non-face-value reading does not so commit. A Stern-style operator analysis of how an utterance like 1 comes to communicate its non-face-value content might go something like this:

1*. She has Mthat ['got a lot of smarts']

where the operator Mthat ['got a lot of smarts'] contributes the set of properties m-associated with 'having the smarts', namely, being intelligent, being quick to understand, and/or having common sense. As with metaphors generally, we should not expect there to be one specific property, and the context will determine which property out of the m-associated set of properties is picked out. For example, in the context of talking about a student who has come top of her class on a competitive philosophy course, the proposition communicated will be something along the lines of "She is very academic". In the context of talking about a child successfully persuading her parents to give her more ice-cream against their considered judgement, the proposition communicated will be something along the lines of "She is good at persuading people to get what she wants". Now, speakers and hearers do not consciously register the operator in the logical form of their uttered sentences, nor do they consciously select the appropriate property or set of properties for the context. Rather, they engage in the same automatic process of meaning selection that they engage in when processing face-value language. That is to say, they directly retrieve the most contextually appropriate content, be it face-value or non-face-value.

The content fictionalist who wishes to go down the metaphor fictionalist route, then, might revise the story operators to resemble metaphor operators. On this view, the modal sentence:

2. There could have been blue swans

would read:

2*. There Mthat['could have been'] blue swans

where the operator Mthat ['could have been'] contributes the set of properties m-associated with 'could have been', perhaps: being such that according to possible world modal realism, there is a world where there are blue swans; being such that a few changes to the actual world would result in blue swans; etc. Crucially, as in the case of 'smarts', speakers and hearers need not consciously register the operator, the m-associated properties (here, where the story operator makes an

appearance) or consciously choose the correct reading in context. Rather, speakers and hearers directly apprehend the non-face-value reading.

The same goes for metaphorical interpretations of sentences of the form:

3. S knows that P

which would be interpreted as having the structure:

3* S Mthat ['knows'] that p

where the m-associated properties contributed by ['knows'] include that S is a good informant with regards to P. (or alternatively: that S is justified in believing that P, that S has good evidence for P, etc.)⁷².

This version of operator fictionalism explicitly builds the operator into the underlying structure of the sentence. The operator has the role of being the mechanism by which uttering the sentence comes to communicate the semantic content of the utterance, specifically, by contributing the m-associated set of properties that furnish the semantic content of the sentence. On this version of operator fictionalism, we should therefore not expect the operator to be transparent to speakers and hearers, for it is a mechanism embedded in the grammatical form of the sentence. The phenomenological objection thus loses its bite: we should not expect speakers to be able, via introspection, to work out whether their uttered sentences contain operators, and what exactly the content communicated via utterances of sentences that contain operators in their logical form is.

Still, one might worry that this response misses the real force behind the phenomenological objection: the problem is rather that in the discourse domains given fictionalist treatment, speakers take themselves to be communicating the face-value content. Speakers think that they are talking about numbers, the world, knowledge and moral facts. So the problem is not that speakers do not realise that they are communicating a non-face-value content via sentences which contain operators, but rather that speakers do not realise, and perhaps would not recognise, the 'metaphorically associated' properties as the properties that they wish to communicate about. In typical cases of metaphor, the reverse seems to be the case: speakers take themselves to be communicating the non-face-value content; it is just that it is not always apparent to them that their utterances, at face-value, mean something else. That is to say, speakers who utter the metaphorical sentence "We're back to square one" would recognise the m-associated properties of 'back to square one' as the properties that they wish to communicate, namely, being at a point in a project where the progress that one had taken themselves to have made is void, and one must start again. Speakers might not recognise, without prompting, that the face-value reading is something akin to 'physically returning to a square, on the background assumption that progress corresponds to moving along numbered squares'. In contrast, speakers who utter "S know that

⁷² Stone's (2000) operator view of knowledge discussed in §3.3.2. shares many similarities with this view. But because Stone does not make use of the notion of fiction in constructing his operator view, his view is not best described as fictionalist.

P” might not recognise the non-face-value content as what they are aiming to communicate (“S is a good informant with regards to P”) but might recognise the face-value content (“S is infallible with regards to P”) as what they are trying to communicate. So, an objection goes, there is a disanalogy between typical metaphors and the fictionalist’s target discourse, such that the latter still falls prey to the phenomenological objection.

In response, it is worth noting that establishing precisely what the face-value content of utterances amount to is not straightforward. Indeed, much of the motivation for fictionalism, including metaphor fictionalism, comes from the realisation that the face-value content is either indeterminate (Yablo 2000, Chung 2019) or ontologically commits speakers to entities that they would not wish to commit to (Yablo 2000, Joyce 2005). This should put some pressure on the claim that speakers take themselves to be communicating the face-value content of their utterances. Although speakers might readily agree that they are 'really' talking about numbers, modality and knowledge, few could give a detailed and accurate paraphrase of what they mean by the terms 'numbers', 'modality' and 'knowledge'. Given this, we should perhaps be sceptical of the claim that speakers intend to communicate the face-value content of their utterance. More plausibly, speakers do not realise that the content which they intend to assert (and which they believe) *is* the content communicated by a non-face-value reading of their utterance. But as we have seen, this is compatible with a Stern-style metaphor fictionalism, on which the non-face-value reading is directly parsed as the semantic content of the uttered metaphorical sentence.

Second, one might question the supposed disanalogy in the other direction. That is, one might question the extent to which speakers and hearers engaged in talk that involves common and everyday metaphors consciously communicate the non-face-value content. Returning to Yablo's example 1 above - 'she's got the smarts' - speakers and hearers indubitably see, or can be made to see, that they are talking about intelligence (or one of the other set of properties m-associated with 'the smarts'). But this is not to say that they do not also take themselves to be talking about 'smarts'. Yablo's insight is that 'smarts' is ontologically committing in ways speakers would reject, but the ontological commitment is not perhaps salient or apparent to most speakers. Speakers and hearers directly retrieve the non-face-value content from utterances of 1, but compatibly with this, they also take themselves to be making claims about 'smarts' that amount to the same thing. Given this, the metaphor fictionalist need not accept that there is a significant disanalogy between how the fictionalist treats its target discourse, and how metaphor theorists treat metaphorical utterances. Insofar as metaphor operators do not fall foul of phenomenological objections, fictionalism that turns on similar operators also does not.

Finally, it is worth noting that not all metaphor theories on which the non-face-value content is directly communicated posit tacit operators. Wearing's version of content metaphor theory drops the operator, and instead holds that metaphorical sentences directly communicate their non-face-value content in virtue of category flexibility. In particular, Wearing argues that what it takes to count as an extension of the relevant category depends on features of the context, so that uttering the sentence “Juliet is the sun” asserts the proposition that ‘Juliet is the sun’, but what it takes to count as ‘the sun’ varies according to context. Hearers pick out the most appropriate interpretation of the category 'sun' given its context. To use Wearing's own example, the metaphorical sentence 'the woods are laughing' invites hearers to understand "how the woods

count as laughing in the context in which they are so described" (2006:324). The non-face-value content of the unit 'is laughing' is whatever the woods are doing to count as laughing in the context, for example, making a particular sound as the breeze drifts through the leaves. Although this is not what would typically count as laughing in most contexts, it is nevertheless 'an attenuated sense of laughing'. Hearers draw upon a variety of factors in picking out what it takes to 'count as an X' in the given context, including features of the conversation in which the metaphor is being used, facts about the portions of the world that the metaphor is being called on to describe, and the extent to which Y comes close to really being an X given the present purposes. In the laughing woods example, the sound that the woods are making count as laughing if, given the current purposes, it is sufficiently similar to the sound of real laughter⁷³.

So, positing operators is not the only option for the content metaphor fictionalist. Adopting a Wearing style of metaphor theory would yield a fictionalist theory on which what counts as 'knowledge' in a context is dependent on a range of factors, including the conversational context in which the sentence is uttered and how closely the epistemic state under consideration resembles knowledge given the current contexts and purposes. As was the case for the operator version of metaphor fictionalism, this type of metaphor theory is well equipped to handle the phenomenological objection, for this view also relegates the redirection part of content fictionalism to a linguistic mechanism which speakers and hearers are unaware of. On this view, speakers and hearers unconsciously use the appropriate category for the context and read the non-face-value content straight off the sentence, even when this means that the category is 'metaphorical'. Therefore, it is no criticism that speakers take themselves to be talking at face-value and to be asserting a face-value content.

5.1.6. Metaphor Fictionalism: knowledge talk shouldn't feel like fictional talk

To conclude the discussion on the phenomenological objection, I have argued that metaphor fictionalism is particularly well suited to answering the phenomenological objection to fictionalism, for metaphors are an example of non-face-value talk that feels nothing like engaging in fiction, where 'engaging in fiction' is limited to playing games or watching shows. First, I have argued that there are interesting and fruitful parallels in the debate between force and content fictionalists on the one hand, and the debate in metaphor theory concerning whether metaphorical sentences are assertoric, on the other. I have argued that although a fictionalism underpinned by a pretence view of metaphor is not well-equipped for handling the phenomenological objection, the metaphor fictionalist has other options. The force fictionalist can adopt an implicature metaphor theory as the basis for her fictionalism, thus avoiding the objection that fictionalism disrespects speakers' phenomenology as regards their attitudes and speech acts. And the content fictionalist can adopt an operator metaphor theory as the basis for her fictionalism, thus avoiding the objection that fictionalism disrespects speakers' phenomenology as regards the content of their assertions. The simple answer to why knowledge

⁷³ Here is how Wearing (2006:325) puts the point: "an object ' counts as an X ' when, by making appropriate presuppositions and hypotheses and imaginings, taking into account what we know about our words and the world and what we're talking about at the instant, we can see how its features could be features of X ' s, and its behaviour X-ish behaviour".

talk doesn't feel like fictional talk, then, is that metaphor fictionalism predicts that knowledge talk *shouldn't* feel like a fiction.

5.2. Dead Metaphors

So far I have distinguished between metaphors that are saliently metaphorical (“Juliet is the sun”) and metaphors that do not obviously strike hearers as metaphorical on first hearing, perhaps because they are so familiar (“I’m back at square one”). However, there is a deeper concern regarding familiar metaphors which goes beyond phenomenology, and which I have not so far addressed. The worry is that these metaphors are not merely familiar, they are *dead* (or *frozen*): that is to say, their use to communicate the non-face-value content is so familiar that they have ceased being metaphors altogether. The idea is that familiar and conventional metaphors like “go back to square one”, “mouth of the river”, “feel down” have over time become literal synonyms for their communicated content. When these dead metaphors first came into circulation they were live and processed as metaphors, however with time they have come to be processed as literal ways of asserting their communicated content. The worry is that even if sentences of the form “S knows that P” are metaphorical in some sense, they are more akin to ‘mouth of the river’ than to live metaphors. If knowledge talk consists in dead metaphors and dead metaphors are not properly considered metaphors - and crucially, if they do not function as metaphors - then knowledge talk is not metaphorical and does not function as a metaphor⁷⁴.

What, then, kills a metaphor? Davidson (1978) thinks that a metaphor dies when it ceases to prompt hearers into noticing a similarity between the source and target domains. He considers the expression ‘bottle mouth’, and theorises that in the past this expression caused hearers to notice a similarity between animal mouths and the opening of bottles. With time and increased exposure, we have stopped noticing a similarity to animal mouths upon hearing the expression “bottle mouth”, and so ‘mouth’ has evolved into a literal word for the opening of a bottle. Compare this to a poetic metaphor which uses the word ‘mouth’: “the city sleeps with death upon her mouth having a song in her eyes” (from E.E. Cummings’ 1923 Impressions IV). On a Davidsonian picture of metaphor this sentence prompts speakers to notice the similarity between a woman sleeping with the traces of something that she has eaten on her mouth, and the city becoming quiet as its inhabitants retire for the night after a hard-worked day. Compare it also to the everyday, but still noticeably metaphorical expression “the experience left a bad taste in the mouth”. According to Davidson hearers are directed towards understanding that the unpleasantness of the experience continued to negatively affect the speakers’ memory and impressions of it even after the experience had finished, and are so directed by being prompted into noticing the similarity between experiences of this type and bad tastes which linger even after the unpleasant food has been eaten. In contrast to both Cummings’ highly poetic metaphor and the more conventional one, ‘bottle mouth’ and ‘river’ mouth do not bring to mind literal mouths, let alone encourage noticing any similarities.

⁷⁴ Without making reference to metaphors, Hawthorne expresses the worry that even granting that sentences of the form “S knows that P” started off with a face-value content that is different to what they now convey, “over time, the propositions conveyed would be fossilized into the semantics” (2004, chapter 3, footnote 22).

Davidson's criterion for establishing whether a metaphor is dead or alive strikes me as overly stringent. For Davidson seems to require that metaphors actually, typically, induce hearers to notice similarities. Yet, many metaphors might be 'resurrected' so that the similarities become noticeable, even if they do not typically cause us to notice similarities. For example, modifying ("that horse you're flogging isn't dead yet, but it sure isn't well"), extending ("She's got the smarts, and the geniuses too") or alluding to a conventional metaphor can make otherwise unnoticed similarities salient, suggesting that these metaphors are not altogether dead. Joyce (2020) proposes a modification to Davidson's view on which a live metaphor need not actually induce us to notice similarities, but less demandingly that they in principle be capable of prompting such similarities. As such, "all those dead metaphors rise from their graves, for even something as dull and familiar as "the mouth of the river" can, with a small effort, be used to conjure a comparative image" (Joyce, 2020:114). Joyce contrasts this to "salinity of the river" which he claims cannot evoke comparative images.

First, I think that Joyce underestimates the extent to which similarity can be found even in typically literal terms like 'salinity'. Consider the sentence "How much salt water thrown away in waste / to season love that of it doth not taste" uttered by Friar Laurence to chide Romeo for despairing over his unrequited love for Rosaline, only to quickly turn his affection to Juliet. 'Salt water' is literal - tears are literally salt water - but its use in the metaphorical sentence overall communicates that Friar Laurence takes Romeo's tears to be performative. Second, Davidson frames metaphor death in terms of actually noticing similarity (rather than potentially noticing similarity) because, as he frequently remarks, "everything is like everything" (1978:39). Any two domains are similar along some dimension or other. A metaphor is apt if it generates relevant and illuminating comparisons which enable hearers to think about the target domain in new and fruitful ways. But in principle, comparison of one type or another can be generated between any two domains. Given sufficient context, even entirely literal sentences can be used to generate a comparison. For example, Reimer (2001) argues that comments about other peoples' appearance (e.g. "those two are in great shape") can prompt hearers to notice, and draw comparisons to, their own physical appearance - this is so in spite of the sentence being entirely literal. Whereas Davidson's criterion is too stringent, Joyce's criterion is over-generous, counting many literal utterances and dead metaphors as live metaphors.

To my mind, where both of these accounts go wrong is in allowing phenomenology (or in Joyce's case, potential phenomenology) to demarcate whether a metaphor is dead. Rather, taking a hint from Walton, I want to suggest that a metaphor is dead if the source domain has ceased to shape or influence how we think about the target domain. Walton writes that a metaphor is live if it's "implication or introduction or reminder of a game of make-believe" (1993:47) continues, or if they "invoke, remind us of" the game of make-believe. Walton's notion of implication and introduction is not phenomenological - it does not refer to whether speakers and hearers recognise that they are taking part in a game of make-believe, but rather to whether a game of make-believe is in play. Crucially, the source and target domain need not be connected - invoked, reminded of - by similarity, but rather are connected in virtue of the way in which the source domain structures talk about the target domain. Walton naturally situates this view of dead metaphors within his particular account of metaphor, on which frames are games of make-believe

constructed from principles of generation which determine the content and truth value of the non-face-value assertion. But we can recast his crucial point in more general (less committing) talk of frames, and say that a metaphor is dead if the target domain is no longer structured by the source domain.

A possible example: the term 'drive' as in 'to drive a car' originally, literally, meant 'pushing from behind' and applied chiefly to the act of herding livestock by standing behind them and prompting them to move. We can imagine that when the term first came to be applied to the act of moving vehicles, it was used metaphorically. Perhaps the frame 'driving [animals]' partially structured how people thought about the act of moving vehicles. Today, the term 'drive [vehicles]' is entirely literal, and has as a face-value content "to move vehicles". But crucially, the reason that it is now literal (if it ever was metaphorical) is not that we have ceased to notice a similarity between herding livestock and moving vehicles, but rather because the source domain of herding livestock no longer structures how we think about moving vehicles. If 'drive a car' ever was metaphorical, it is now a dead metaphor. In contrast, the metaphor for crime as a beast (as in, 'crime stalks the city') shapes how crime and those committing it are thought about, and so is rightly considered a live metaphor.

This suggestion echoes in part remarks made by Lakoff (1987) who complains that the term 'dead metaphor' fails to discriminate historical metaphors from conventional metaphors. Historical metaphors are metaphors like 'drive' and his own example of 'pedigree' (from the French word for 'crane's foot', which family trees were thought to resemble) in which the target domain no longer maps onto the source domain at all, because the structure or even the linguistic mapping is absent. In contrast, conventional metaphors are ones that are widespread but in which the target domain continues to be mapped in terms of the source domain. Lakoff's examples include 'grasp' as in 'grasp an idea'; for Lakoff, the mapping UNDERSTANDING IS GRASPING in which ideas are treated, metaphorically, as physical objects in space which can be held underpins how we think about understanding. Importantly, we need not buy into Lakoff's expansive conception of metaphor to fruitfully draw the distinction he suggests between historical metaphors and conventional metaphors. The crucial point is that a metaphor is dead when the source domain ceases to structure the target domain, and therefore whether a metaphor is dead might not be readily discernible from introspection.

This way of construing what it means for a metaphor to be dead has several advantages. It follows Camp's recommendation to see dead metaphors as falling at "one extreme on a spectrum from conventionality to novelty" (2006:290), with most everyday metaphors lying somewhere in between. This contrasts with theorists (e.g. Wearing 2006) who view dead metaphors as a distinct phenomenon altogether. The extent to which a source domain structures a target domain is a matter of degree, and so the extent to which a metaphor is dead or alive is likewise a matter of degree. Second, this view easily answers the question of how a dead metaphor acquires its literal content. Reimer (1996) notes how on Davidson's view, a metaphor communicates only its face-value content (e.g. that my surgeon is a literal butcher), but once dead it acquires a different face-value content (e.g. that my surgeon is careless) - but there is no explanation for how the sentence comes to suddenly acquire a very different meaning to the one that it had previously. On the structuring view, the sentence possesses the non-face-value content that it does in virtue of the

way in which the source domain structures it. Upon death, the metaphor acquires as a face-value content what previously was a non-face-value content. Finally, this view remains neutral on the question of whether a metaphor communicates a non-face value content, implicates a non-face value content, or merely calls attention to some state of affairs. On this picture, a metaphor has died if the target domain is no longer structured in terms of its source domain, irrespective of whether the structure is at the level of asserted content.

Returning to knowledge, the knowledge metaphor is alive if what we mean by alive is that the source domain continues to shape and structure how we think about the target domain. The knowledge metaphor does not strike us as metaphorical, to be sure; but this is attributed to the knowledge metaphor being a conventional (and therefore, unnoticed) metaphor, rather than to it being a dead (and therefore, non-structuring) metaphor.

5.3. What fixes the content of typically non-face-value talk?

It might seem as if theories which divorce face-value content from what is conveyed have a difficult time explaining how the face-value semantic content of the target discourse gets fixed in the first place. Defenders of both force and content epistemic fictionalism hold that sentences of the form “S knows that P” are typically false if taken to assert the content communicated by a face-value reading, but crucially, the content communicated by a face-value reading is never (or is very rarely) what these sentences are used to communicate. But this raises the question of how sentences of the form “S knows that P” acquire their face-value meaning to begin with, despite the fact that standard usage of these sentences conveys some other meaning altogether. Why has the non-face-value meaning not simply become the face-value meaning, to be directly communicated through assertion?

Here is how Blome-Tillman (2013:4303) puts the point:

[i]f we adopt the eminently plausible view that the linguistic meaning of an expression is determined by its use in a speech community—or by the conventions governing its use—then there seems to be no room for a difference between literal meaning and standard speaker meaning.

And here is Hawthorne (2004:143):

And if the semantic value of ‘know’ rarely if ever plays a role in the propositions that are conveyed and believed when knowledge ascriptions get used, then where does its purported semantic value come from? The suspicion remains that [pragmatic views] posit too great a gap between semantic value and linguistic practice.

So, according to the plausible assumption that face-value meaning is fixed by ordinary use, sentences of the form “S knows that P” are typically used to communicate their face-value meaning. The epistemic fictionalist will therefore struggle to explain the posited gap between face-value and non-face-value usage.

There are several replies available to the epistemic fictionalist. Dinges (2016) argues that we are not forced to accept what Blome-Tillman describes as the “eminently plausible view” on semantic content acquisition, and instead (taking a hint from the debate on loose talk), that we should modify the view to say that semantic content is determined by use in a speech community *when members of that community are using the term literally*. For example, a loose talk analysis of the sentences “3 oc’clock” and “this line is taking forever” would hold that these sentences are not typically used to communicate their face-value content; but nevertheless, they are used literally on some very rare occasions. It is their meaning on these rare occasions that determines their semantic content, that is, that determines the content that such sentences communicate on a face-value reading. Likewise, Dinges suggests that “on the (potentially rare) occasions where “knows” is used literally (for example, when we speak of a priori truths or immediate perceptions), it does convey what it literally communicates; namely, a quite demanding epistemic relation” (2016:2585). It should be noted that the view that Dinges discusses in relation to this objection holds that the face-value reading of knowledge is a demanding, but not infallible, epistemic relation. One might worry that if knowledge is given the face-value reading posited by the sceptical infallibilist, then it becomes less plausible to claim that there are occasions where the face-value content is the content that the speaker intends to communicate.

Less plausible, perhaps, but not impossible. Chung (2018) argues that Hawthorne’s way of framing the problem assumes, without warrant, that standard usage in co-operative contexts is what matters. If we expand the ‘standard usage’ datum to include antagonistic contexts then we see far more use of potential face-value use of “S knows that P” - and, crucially, of “S doesn’t know that P”. Not only are philosophical sceptical challenges one such example, but everyday challenges such as “the press making a mistake, a person lying, or an unforeseen disaster occurring are frequently admitted to be knowledge-imperilling when they are raised” (Chung 2018:406). If we include challenges into the pool of standard usage for knowledge talk, then, it becomes more plausible to claim that there are face-value uses of knowledge which determine the semantic content of sentences of the form “S knows that P”. Furthermore, even granting that literal talk of infallibility is now not frequently encountered in everyday talk, this was plausibly not always the case. To give an easy example, any culture with a religion featuring an omniscient God would have had reasonably frequent reason to think and utter sentences of the form “God knows that P” where the communicated content is the content expressed by a face-value reading of the sentence, and where this is also the content posited by the infallibilist. On the additional assumption that such talk has been culturally significant enough to seep into how we today use the concept, then historically frequent usage might also be included into the datum used to determine the face-value reading of a concept.

Although I think that Dinges and Chung are essentially right to say that even a very (perhaps, maximally) demanding knowledge concept may be used at face-value as a standard (albeit, rare) usage, an epistemic fictionalism informed by the function-first methodology need not argue this point. The function-first-fictionalist has another response at her disposal. For on function-first methodology, as we saw in Chapter 2 of this thesis, concepts come equipped with the meaning that makes them functional. That is to say, it is a concept’s function, rather than its standard use, which we must look to in determining the contours of the concept. And although some special

explanation is required in these cases, a concept might nevertheless come equipped with a face-value meaning that is divorced from what a sentences employing that concept are typically used to convey. Indeed, this possibility is already hinted at in a few suggestive comments by Chung (2018:407), who writes that:

Some uses of an expression might be more central to the achievement of our aims than others, and because of this, one might think that those uses should be accorded greater importance as far as determining its semantic content is concerned. [...] in order to properly answer a question of the form “What is the semantic content of the expression *e*?” we would first need to be able to answer the question “What is the (primary) purpose of the expression *e*?” In order to ascertain the semantic content of “knows,” then, we would first need to provide an answer to the question: What is the (primary) purpose of the expression “knows”? Or, in other words, what is knowledge-talk (principally) for?

Chung goes on to suggest that the role that knowledge talk plays in co-ordinating inquiry is what fixes its face-value usage. In a familiar Craigian move, Chung writes that whereas the ‘ending inquiry’ function might be expected to push the standards for knowledge down, the ‘beginning inquiry’ function pushes them up towards much more demanding standards – perhaps, even, to the maximally demanding standards posited by the infallibilist. Reasoning that to end inquiry one must first begin it, Chung submits that the latter is the more fundamental use, and therefore the one more central both in achieving the aims that the concept of knowledge exists to help us achieve, but also, crucially, in determining the face-value content of sentences of the form “S knows that P”.

Chung’s principal point is that the fictionalist has access to avenues for answering the question of how sentences of the form “S knows that P” acquires its face-value content, and that these avenues need not rely on intuitions regarding standard use. The key insight is when we consider purpose to be an essential role, then it is no longer true that “the semantic value of ‘know’ rarely if ever plays a role in the propositions that are conveyed and believed when knowledge ascriptions get used”. The epistemic fictionalist view that I have developed goes beyond Chung’s suggestion to vindicate not only that knowledge might have the face-value content posited by the infallibilist, but crucially that in order to be maximally functional, the concept of knowledge needs to be structured precisely so that there is in fact a “difference between literal meaning and standard speaker meaning”. That is to say, not only is there a purpose served by knowledge coming equipped with maximally demanding standards at the face-value level; but there is actually a purpose served by knowledge coming equipped with maximally demanding standards at the face-value level, while coming equipped with different standards at the non-face-value level.

I argued in Chapter 4 that knowledge talk being metaphorical functions to produce a concept that communicates that one has cleared the threshold for being a good informant for members of their community. On the one hand, being framed along the lines of infallibilism makes salient the demanding standards needed for being a reliable informant. On the other hand, it enables agents to treat as infallible pieces of information that have already been tagged as sufficiently

reliable. Thus on this view (as with metaphors more generally) the face-value infallibilist meaning of the concept of knowledge does play a role towards what proposition conveyed by sentences that include it. In particular, it plays a framing role. Its role is to frame and structure the face-value meaning that is conveyed, and to do so in a way that speaks to the function that the concept of knowledge fulfils. Therefore, it acquires the face-value meaning needed to frame a concept of knowledge capable of fulfilling its various functions.

5.4. Conclusion

This chapter has discussed and tried to answer what I consider to be the three most pressing objections to the metaphor-based epistemic fictionalism that I outlined in the last chapter. These objections target the view at different levels of analysis: the phenomenological objection is typically levelled at fictionalism in general; the objection that widespread metaphors are dead metaphors targets only metaphor fictionalism; and the objection that fictionalism struggles to explain where its target discourse acquires its face-value semantic content targets not only all fictionalist views, but many other non-fictionalist pragmatic views also. I have tried to show that all three objections not only can be answered, but that they can be answered in ways that further lend support to the function-first epistemic fictionalism that I have developed. Along the way, I have also attempted to demonstrate how both a content and force version of epistemic fictionalism could be developed, and that whether one's fictionalism ends up as a force or content fictionalism will depend on which theory of metaphor it employs. This is not a question that the epistemic fictionalist needs to settle – rather, I have argued that choice of underlying metaphor theory should be guided by which metaphor gives the best account of metaphor in general, rather than the best account of fictionalism. For the epistemic fictionalist, what matters is that there are plausible metaphor theories that cohere well with epistemic fictionalism.

To briefly summarise, the phenomenological objection is generally considered a devastating objection to fictionalism. I have argued that a fictionalism ground in metaphor theory (where 'metaphor theory' is either a non-pretence force view of metaphor or an operator-based content view of metaphor) easily explains why knowledge talk doesn't induce the phenomenology of engaging in fiction: using metaphorical talk is phenomenologically similar to using face-value talk. Further, this is the case on at least two plausible views of metaphor, which yield a force and content epistemic fictionalist theory respectively. Therefore, I hope that this solution might be of interest not only in the context of epistemic fictionalism in particular, but to fictionalists more generally. In response to the question of dead metaphors, I have argued that so long as the source domain continues to shape and structure the target domain by shaping and structuring talk, thought, and practices concerning the concept, then the knowledge metaphor should be regarded as alive. In response to the final objection about face-value content acquisition, I have employed a function-first lens to argue that reflecting on the purpose of the concept of knowledge opens up the possibility that a metaphorical knowledge concept would acquire the face-value meaning that would usefully shape and structure the way that knowledge talk is used. This is the case even if there is a gap between the face-value meaning and the non-face-value meaning. In doing so, I hope to have further made the case for how a fictionalist can fruitfully apply the function-first methodology to analyse her target non-literal talk. Having addressed what I take to be the most

pressing objections for epistemic fictionalism on its own terms, the next chapter compares epistemic fictionalism to its main rival theories.

Chapter 6: Epistemic Fictionalism and its Rivals

This chapter continues the defence of epistemic fictionalism by comparing it to rival theories of knowledge. The first half of the chapter defends the claim that knowledge talk is metaphorical, as opposed to belonging to a related but distinct category of non-face-value speech such as hyperbole and loose talk (§6.1.), or regulative ideals (§6.2.). I argue that among other advantages, epistemic fictionalism better captures the function of knowledge talk. Then, I compare epistemic fictionalism to what I consider to be as its main rivals: epistemic contextualism (§6.3.) and sceptical pragmatic invariantism (§6.4.). Both of these views also posit that sceptical infallibilism is in an important sense right about (at least one of) the meaning(s) of sentences of the form “S knows that P”. These views appear at first pass to generate the same results as epistemic fictionalism, without the burden of positing non-face-value mechanisms. In response, I argue that there are ‘counterpart’ ways of specifying epistemic fictionalism for both epistemic contextualism and sceptical pragmatic invariantism, and that these counterparts retain the advantages of epistemic contextualism and sceptical pragmatic invariantism respectively, while additionally resolving objections to the face-value version of each view.

6.1 Hyperbole and Loose Talk

The account that I’ve laid out has it that the community-indexed relevant possibilities are metaphorically framed as *all* of the possibilities. This might generate the worry that the pragmatic mechanism underpinning the redirection from face-value content to communicated content is not metaphor but hyperbole. Hyperbole (or exaggeration) is a widely recognised form of non-literal discourse in which a false face-value content is used to communicate a true communicated content. These sentences are typical examples of hyperbole:

- (a) I have been on hold *forever* (when I’ve been waiting for an hour)
- (b) This job would be *perfect* for me (when this job would suit me very well)
- (c) My slice of cake is *tiny* (said by a 6 year old boy whose older sister’s slice is only slightly larger)⁷⁵

On the standard reading, hyperbole adds emphasis and effect to the communicated content by highlighting the contrast between the face-value content and the actual situation (see Colston 1997, Colston and O’Brien 2000, Colston and Keller 1998, McCarthy and Carte, Rubio-Fernandez et. al.). As McCarthy and Carte (2004:158) explain, “hyperbole magnifies and upscales reality, and, naturally, upscaling produces a contrast with reality”. In sentence (a) above, the hyperbolic “forever” emphasises that an hour is a long time to wait by calling the hearer’s attention to the contrast between waiting forever and waiting for an hour. One might worry that the account I have described suits a hyperbolic account better than a metaphorical one, with the end of the scale (“all of the alternatives”) being the face-value content of sentences of the form “S knows that P”.

⁷⁵ Sentence (c) from Carston and Wearing (2015:80)

That knowledge is a form of hyperbole has been suggested by a number of philosophers, among them Unger (1975), Bonjour (2010), Carter (1995) and Schaffer (2004). These philosophers employ a somewhat unorthodox notion of hyperbole, on which it is a form of loose talk. Unlike hyperbole, loose talk typically succeeds in communicating the non-face-value content by implying that for current purposes there is no difference between the standards governing the face-value content and the situation in question. On a loose talk reading, sentence (a) communicates that there is no practical difference between waiting forever and having waited for the very long time that one has actually waited for. On this view, one's epistemic situation is sufficiently similar to knowledge as to render the differences insignificant for one's practical purposes. For example, Bonjour proposes that sentences of the form "S knows that P" function to indicate that one is in close approximation to knowledge, such that for current purposes to distinguish would be an instance of "fine hair-splitting" (2010:73).

Unger also suggests that sentences of the form "S knows that P" typically indicate that one is close enough to knowing for the purposes at hand. His account differs from the others in that it is error-theoretic: on Unger's view, everyday users of knowledge discourse are unaware that the face-value content of their utterances are false. This is explained by the difference between the face-value content and the communicated content being so insignificant for practical purposes as to go unnoticed. Nevertheless, it is easy to see how a hyperbole account can be derived from Unger's error-theoretic account, as indeed Schaffer (2004) does. Building on Unger's work, Schaffer argues that if taken at face-value, sentences of the form "S knows that P" are false because they would communicate that S has eliminated every possibility in which $\sim P$. However, the face-value content generates a number of entailments, some of which will be true and informative to audiences (e.g. that S has eliminated the relevant possibilities) and some which will be false but not misleading (e.g. that S has eliminated the possibility that she is a BIV). On Schaffer's view, using hyperbole enables speakers to strike the optimal balance between effect and effort, communicating the relevant information with the least time and effort. Sentences of the form "S knows that P" communicate the true and informative entailments, while (due to presumed audience co-operation) stopping short from expressing the false entailments⁷⁶.

Carter is the exception to the trend of framing knowledge as hyperbole in terms of loose talk. For Carter, knowledge talk is hyperbole in the standard sense of emphasising the contrast between one's actual epistemic situation and knowing in order to add flourish and emphasis to the communicated content, in a way that is obvious to hearers. The idea is that speakers say something much stronger to add emphasis to "S guarantees that P". Carter does not specify

⁷⁶ In addition to these, several more theorists advocate or discuss a non-hyperbolic loose talk view of knowledge on which sentences of the form "S knows that P" communicate that one is close enough to knowing for present purposes (see e.g. Davies 2007, Fantl and McGrath 2007, Conee 2005). The difference between loose talk and hyperbole views in the epistemology literature is that loose talk theorists posit a moderate invariant face-value content for knowledge rather than the sceptical one held by hyperbole proponents. This distinction tracks the central position that sceptical arguments enjoy in epistemology, rather than a fundamental difference in the linguistic functioning of loose talk and hyperbole themselves, since hyperbole need not involve maximal standards in the face-value content (e.g. "I have been waiting for a thousand years [as opposed to forever]", "This is the best restaurant in Scotland [as opposed to the world]", "Her house is a mansion [as opposed to a castle]"). I here limit the discussion to views on which sentences of the form "S knows that P" have the face-value content posited by the sceptic.

exactly why our language habits have developed towards emphasising the communicated content, but he does make a suggestive comparison. Carter notes that hyperbole in general is often used for purposes other than to describe, and gives recommendation as his one example: “To say that a film is the best ever seen by a film enthusiast with good taste is to commend the film highly. The function of hyperbole in this case is obviously that of commendation by means of conventional use of exaggeration” (1995:61). Knowledge as recommendation ties in well with the function-first methodology being used throughout this thesis, on which one of the posited functions for knowledge is to recommend reliable informants. We could develop Carter’s account into a view on which the strength of S’s epistemic position is exaggerated to maximal strength in order to recommend S as an informant, just as the goodness of a film is exaggerated to maximal goodness in order to recommend the film as enjoyable. Crucially, in order for hyperbolic sentences to be registered as recommendation rather than a description, hearers must register that the speaker is talking hyperbolically by noticing the contrast between ones situation and the face-value content of the utterance.

It is important not to overstate the differences between hyperbole and metaphor. On some views (e.g. Sperber and Wilson’s relevance theory 1985, 1995, 2008) there is no interesting difference in interpretation or processing between metaphor and hyperbole, and on many views at least some instances of non-literal talk are both metaphorical and hyperbolic (e.g. Carston 2002, Carston and Wearing 2011). Even admitting a genuine distinction between metaphor and hyperbole, both are non-literal tropes in which speakers say one thing to communicate another and as such share many advantages as accounts of knowledge discourse. Both for example would explain the appeal of and resistance to sceptical arguments; both appeal to a non-face-value mechanism to explain the redirection in communicated content towards a non-face-value content; and both enjoy solid precedents, thus avoiding the problem of making knowledge discourse an outlier. In conceding that the sceptic has the correct face-value reading of knowledge sentences, a hyperbole view would still be a radical departure from traditional analyses of knowledge.

Having said this, I do think that metaphor gives the better explanation of the patterns of use we see in knowledge talk. To start, knowledge does not behave like an instance of hyperbole, as demonstrated by it failing a series of tests. These include: (a) it is obvious to speakers that hyperbole is false on a face-value reading or at least, this fact can be made obvious to them. In the case of knowledge talk, it is not the case that speakers register their utterances as hyperbolic (Hawthorne 2004, Blome-Tillmann 2013); (b) hyperboles are typically uttered in a special tone of voice, which involves accenting the specific hyperbolic word and using a wider pitch range for the utterance as a whole (Carston and Wearing 2015), while knowledge claims are typically uttered in ordinary tones (Hawthorne 2003); (c) hyperboles are often combined with other non-literal tropes to add emphasis, while knowledge talk tends to stand alone; and finally (d) hyperbole is easily cancellable by being prefaced with terms like “literally”, while knowledge sentences are not cancellable in this way (Hawthorne 2003, Davis 2007)⁷⁷.

⁷⁷ As Hawthorne points out, ‘literally’ has increasingly come to possess a non-face-value use, as in “I literally died” or “I am literally freezing”. Hawthorne suggests that in these cases ‘literally’ contributes to the

Worse, the function of knowledge does not neatly square with the function of using hyperbole. Hyperbole is typically used for emphasis and flair, to stress a point for rhetorical purposes. Knowledge talk is simply not like this. As DeRose notes, knowledge claims are made even in extremely serious and accuracy-demanding contexts such as court rooms, where hyperbole would be inappropriate. In contrast, because metaphors are cognitive aids and are not heard as false, they remain appropriate even in very serious contexts like court rooms. Indeed, several studies have found that metaphoric language is pervasive in court rooms and legal systems at every level of discourse, from lawyers' cross-examinations (Luchjenbroers and Aldridge 2007) and addresses (Cotterill 1998), through to witness testimony (Newman 1999) and to the deliberation processes of judges and jurors (Winters 2001). Indeed, many of the USA's very laws and legal rights themselves are communicated metaphorically (Smith 2007) and the legal process itself is structured along a series of metaphor frames (Thornberg 1995).

A second (and initially more promising) function of hyperbole is to communicate an evaluative attitude towards the state of affairs that the speaker is being hyperbolic about. In particular, it communicates a discrepancy between the state of affairs and the state of affairs that the speaker expected or wanted. Although hyperbole are more frequently used to communicate negative attitudes (Claridge 2011: 81), they need not. As Carston and Wearing (2015:15) put it, "the full range of positive and negative evaluations across all manner of possible scales is available (smaller / bigger /longer /shorter /harder / easier than I wanted/expected; much worse/better than I was expecting or hoping for)". This function looks closer to what Carter has in mind. If sentences of the form "S knows that P" are used to communicate that S is a good informant with regards to P then it is used to evaluate S's epistemic position and to communicate a positive attitude towards it. However, although there is perhaps a general normative sense in which such sentences are evaluative, they are not evaluative in the comparative sense typical of hyperbole. They do not communicate that S is in a better epistemic position than what the speaker expected, wanted or finds typical. Rather, they more closely resemble a description of states of affairs, conveying that S is in a sufficiently good epistemic position for one's purposes, irrespective of the speaker's attitude towards this fact. This descriptive function is more in line with the functions of metaphors, since one of the functions of metaphors is to describe states of affairs in ways that are more cognitively accessible.

As noted, proponents of the hyperbole view of knowledge tend to have an altogether different function in mind: the function of communicating that for current purposes, the present situation is sufficiently similar to the face-value truth of the utterance as to make differentiating the two needlessly pedantic. This is a function typical of loose talk rather than of hyperbole. Even supposing that hyperbole does serve this function, however, there are problems for this view that make it untenable. First, there needs to be greater clarity on what it means to be "close enough"

hyperbole by itself being used hyperbolically, without precluding its ability to typically cancel an implicature. Blome-Tillman (2009) argues that if a sentence which would typically cancel an implicature is itself in violation of a Gricean maxim, then that sentence is best interpreted as communicating some other content - thus it is determinable from context whether these sentences are cancelling to, or adding to, the previous implicature. I think that this is the right interpretation for hyperbolic uses of 'literally'. In some contexts (e.g. "I literally died") to interpret 'literally' as cancelling a hyperbole would be to interpret the speaker as violating the maxim of truth.

to knowing. Lasersohn's (1999) influential loose talk account posits a "pragmatic halo" surrounding each sentence. Each halo is formed of propositions sufficiently close to the face-value reading as to make the difference between them pragmatically ignorable. As Bach (2005) points out, the problem is that for the sceptic, ordinary knowledge claims are never even close to being true. There are always differences between being in even a good epistemic position and being infallible with regards to P, such as the presence of uneliminated possibilities of error. If the errors are sufficiently distant or irrelevant then one might not be required to rule them out (either because doing so is impossible or because they are not taken seriously) but this is not to say that one's position is the *same* as infallibility, or even that it is close. In fact, there will be differences in attitude and disposition: willingness to change one's mind should new evidence come to light; willingness to acquire new evidence on the topic if doing so comes at no or little cost; willingness to rely on the proposition in a different context; willingness to testify that P even in very high stakes; and so on. There are many practical purposes, such as testifying in court, where the difference between knowing and one's epistemic position does matter but in which if one's evidence is sufficiently good then it is still appropriate to claim knowledge. In contrast, the metaphor account does not require that one's epistemic position be close or the same as knowing for practical purposes. All that is required is that the source domain share a characteristic with the target domain, such that the former illuminatingly presents the latter. To say that knowledge is metaphorically framed is not to say that our current situation is close to knowledge, but rather that the state of knowledge can illuminatingly frame certain aspects of our current situation.

Second, there are linguistic concerns with treating knowledge as loose talk. Sentences in loose talk typically succeed in communicating a non-face-value content precisely because their lexical components are at least sometimes used in sentences that communicate their face-value meaning. Saying that Jessica arrived at the party at 2pm when she arrived at the party at 1.58 is an example of loose talk. It is only effective, however, because sentences including '2pm' are at least sometimes used at face-value. Loose talk communicates that for present purposes there is no important difference between such cases where the sentence is true at face value and the current situation. This contrasts with both hyperbole, which often has as its face-value content a limit end of the scale (e.g. *perfect, forever*) which is rarely if ever is given a face-value reading, and with metaphor frames, which also do regularly take concepts that have no or infrequent face-value use as the source domain (e.g. "You're an *angel*"). Finally, the loose talk account makes the wrong prediction in cases of knowledge denial. As discussed, loose talk functions to indicate that the difference between a face-value reading and one's current situation is ignorable for practical purposes. However, as Dinges (2021) points out, the difference between a face-value reading and itself is, trivially, always pragmatically ignorable. This means that even heard with an ear trained towards loose use, sentences of the form "S doesn't know that P" would encompass its face-value reading. This is a problem because on the sceptical reading the face-value reading will be true. Yet, sentences of the form "S doesn't know that P" are heard as false when S has sufficient evidence. If knowledge talk followed the patterns typical of loose talk, we should hear these sentences as true because we should recognise that the face-value content is one of the possible contents being communicated. In contrast, it takes some effort to hear common metaphors as communicating their face-value reading.

Given these considerations, although the metaphor account shares similarities with both hyperbole and loose talk accounts, I think that overall it is to be preferred to both.

6.2 Knowledge as a Regulative Ideal

The view that knowledge is a regulative ideal is relatively underexplored in the literature. There are however three notable exceptions: Butchvarov (1970); Bonjour (2010), who credits the insight to Butchvarov; and Hannon (2019b), who rejects the view. Both Butchvarov and Bonjour support the claim that knowledge is an ideal by highlighting the functions that the concept of knowledge plays.

A regulative ideal is a concept which although unachievable, serves to regulate our practices in a domain. Typical examples include the concepts of *perfection*, *goodness*, *beauty* and *justice*. More controversial examples include *love*, *circle*, *utopia*, *freedom*, *truth* (Butchvarov 1970). These concepts are regulative ideals (henceforth, simply ‘ideals’) because they denote a goal or a standard which can be approached, but not attained. They represent the end point of a scale in each domain, a scale along which we might wish to progress. Their significance lies not in describing states of the world but in “their utility in guiding our thought and action into fruitful and rewarding directions, wholly not withstanding their unrealistic and visionary character” (Rescher 1987:139). Ideals serve to regulate practices because having an ideal state towards which to aspire directs and improves our pursuit of domain-relative aims.

In their respective monographs on regulative ideals, Nicholas Rescher (1987) and Dorothy Emmet (1992) discuss the function of ideals. I think that we can draw out similarities in their accounts, and identify three distinct but related functions which they attribute to ideals. First, **DIRECTION**: ideals orient and direct our actions and practices. Rescher (1987:57) writes that “the key role of an ideal is to serve as an instrument of decision making - a sort of navigation instrument for use in the pursuit of the good”. Ideals function as a tool of decision making by structuring the values, virtues, habits, policies and standards that govern the target practice in general; as Emmet puts it, they furnish “an orientation to our particular evaluations while preserving openness to criticism” (1992: 69). But they also guide individual decisions by enabling agents to set the ideal as a goal, and work towards it. For example, to see the concept *justice* as a regulative ideal is to understand a relevant subset of our practices as being structured in a way that enables progress towards the ideal of justice in that domain. Framing everyday actions in terms of their relationship to the ideal state of *justice* is useful because constructing habits and policies based on those that would advance us towards the ideal of *justice* improves the degree to which our non-ideal just actions approach the ideal (**DIRECTION**).

Second, **YARDSTICK**. Both Rescher and Emmet emphasise the role that ideals play in appraisal that is geared towards taking action. Rescher writes that an ideal “exerts a powerful organising and motivating influence on our thinking, providing at once a standard of appraisal and a stimulus to action” (1987:119), and Emmet that it does so by “giving direction and destroying complacent acceptance of compromises” (1992:59). Ideals, then, afford us with a powerful tool with which to evaluate our current state of affairs, shining a light on what would need to change to bring our current situation closer to the ideal, and motivating us to make these changes. The ideal of a

Utopia as used in political theory, for example, does not aim to be a realistic plan for the future - rather, it functions to draw attention to the features of our present society that are somehow lacking and in need of improvement.

Finally, and relatedly, EFFECTS: ideals play an important psychological role by bolstering commitment and enthusiasm for our practices. Ideals “goad to effort, motivate rather than constrain, urge rather than demand” (Rescher 1987:137). Ideals are valued for their own sake, and attaining them is considered a worthwhile goal in itself. Although attaining ideals is not possible, fixing our sights on achieving an ideal motivates us to make progress towards the ideal. Further, the progress made is valued precisely because it is viewed as progress towards the valued ideal. Returning to the example of justice, conceptualising ourselves as moving towards justice bolsters our enthusiasm for and commitment towards following and enacting the habits and policies that progress us towards justice (EFFECTS). Crucially, “what counts with an ideal is not the question of its attainment but the question of the benefits that accrue from its pursuit” (Rescher 1987:137) - whether an ideal is functional depends upon the functionality of *attempting* to attain the ideal.

Butchvarov argues that the concept of knowledge is also a regulative ideal. In a nutshell, the concept of knowledge denotes an ideal epistemic state towards which it is useful or desirable to aim. This ideal state is as the sceptic says it is - it is a state in which S is infallible with regards to P, in which S's evidence entails the truth of P, in which S has ruled out every error possibility to P, and so on. As Butchvarov says, “it is absolutely impossible that I am mistaken in believing that P” (Butchvarov 1970:59). Therefore, we never (or very rarely) achieve the standards demanded by knowledge. However, the ideal of knowledge nevertheless serves important functions. Indeed, the functions posited by Butchvarov are precisely those that Rescher and Emmet argue that regulative ideals in general serve⁷⁸.

First, our epistemic practices are directed and oriented towards the ideal of knowledge (DIRECTION). As Butchvarov (1970:51) puts the point, “even if [knowledge] were entirely unobtainable, it is precisely that which such [epistemic] pursuit necessarily seeks to approximate. It is that which gives direction to our cognitive enterprise”. Indeed, many epistemologists consider knowledge to be at least one of the principal epistemic goods to which our epistemic practices are geared towards⁷⁹. Orienting the practice of inquiry (to name one principal epistemic practice) towards the ideal of knowledge cultivates epistemic virtues and habits, the possessing of which advances our epistemic standing along the scale towards knowledge. For example, aspiring towards the ideal of knowledge might promote the virtues of intellectual humility and curiosity, and the habits of double checking ones' sources for trustworthiness and ensuring that one has

⁷⁸ Both Butchvarov and Bonjour switch between describing knowledge as exaggeration and ideal. Ideals and exaggeration play a specific role in their arguments: to explain why the maximal standards for knowledge typically go unnoticed. Bonjour is explicit in offering several possible explanations. In light of Butchvarov's comments on the function of knowledge, he is in my view most charitably read as offering an ideals account.

⁷⁹ See e.g. Williamson (2000), Adler (2002), Huemer (2007) and Littlejohn (2013).

sufficient justification in support of one's claims to know. These virtues and habits will aid agents in improving their epistemic position towards knowledge, even if they do not achieve the ideal.

Second, **YARDSTICK**: the ideal of knowledge is the "yardstick with which we can measure the relative adequacy of beliefs" (Butchvarov 1970:61). Considering the gap between the ideal knowledge and our current epistemic state enables us to critically assess our epistemic standing and specify the ways in which it falls short of the ideal, remedying these where it is in our power to do so. On the flip side, measuring the distance between our current position and knowledge might in many instances highlight the strength of our epistemic position, for example by illustrating how the gap is sufficiently small given the aims and stakes operative in the context.

Finally, treating knowledge as an ideal yields the motivational advantages typical of ideals in general, and explains why knowledge is valued (**EFFECTS**). BonJour (2010) argues that theories which respect the intuitions regarding variability of the truth conditions for sentences of the form "S knows that P" cannot also accommodate the intuition that knowledge is valuable. This is because there is no interests-unrelated, principled reason for holding one's current epistemic state as valuable, even when it counts as knowledge for present purposes, if that state could be improved upon and if it would fail to be knowledge in a different set of circumstances. Yet understanding knowledge as an ideal makes sense of why knowledge is valued. As is the case for other ideals such as *beauty* and *justice*, the ideal is valuable in itself, and approximations to it are valuable in virtue of the degree to which they approximate the ideal.

Furthermore, the ideals theorist could posit a similar story to that told about my own metaphor view to explain what sentences of the form "S knows that P" typically communicate. She could posit that utterances of sentences of the form "S knows that P" are not typically used to communicate the proposition expressed by a face-value reading of "S knows that P", but rather that "we apply the concept of knowledge to what in reality are only cases of rational belief because such cases are regarded as sufficiently like knowledge to be, for practical purposes, counted as such" (Butchvarov 1970:59). Sentences of the form "S knows that P" could on this view communicate something along the lines of: *S is sufficiently close to the ideal of knowledge with regards to P for current purposes.*

I am overall sympathetic to the view that knowledge functions as a regulative ideal. Indeed, of the views considered in this chapter, it is the one that most resembles the metaphor view. Both views appeal to a non-face-value mechanisms that play a pivotal role in structuring our practices involving knowledge talk. Whereas the ideals account emphasises the aspirational aspect of knowledge, the metaphor account parses the relationship between knowledge (on face-value) and knowledge (non-face-value) as a mapping. The principal insight of the ideals view is that knowledge is in a sense aspirational, and that the concept of knowledge might function in virtue of, rather than in spite of, denoting a maximal state when given a face-value reading. I think that the metaphor view shares these advantages. The metaphor view too posits that on a face-value reading, 'knowledge' denotes a maximal standard. It is the maximal standard that provides the source domain which 'good enough' epistemic standards are framed in terms of. Furthermore, the metaphor view can to an extent accommodate the aspirational side of knowledge: one of the salient features of the source domain is its aspirational nature.

I will start with an objection to this view that ultimately I think fails. Hannon (2019) argues that even on the assumption that knowledge is a regulative ideal, scepticism is implausible. As discussed in Chapter 1, scepticism comes in two varieties: 'full-blooded', on which no belief is epistemically better (e.g. more justified) than any other; and 'high standards', on which no belief amounts to knowledge, but on which some epistemic beliefs are epistemically better (e.g. more justified) than others. Hannon argues that neither version of scepticism is plausible, on the assumption that knowledge is a regulative ideal. Starting with full-blooded scepticism, he objects that understanding knowledge as an ideal presupposes that the ideal can be approached at least to some degree. Full-blooded scepticism denies this – since no belief is better than any other, no belief is closer to knowledge, either. This leaves high-standards scepticism. However, once the sceptic has admitted (as the high-standards sceptic must) that some epistemic positions are better than others, why think that the concept of knowledge characterises a maximally good epistemic position? Rather, it seems more plausible that the concept of knowledge demarcates one of these other good, but non-maximal, epistemic positions. Further, high-standards scepticism makes it unclear why the sceptical conclusion comes as a revelation, and why scepticism is considered a threat. One would expect speakers to be fully aware that they are progressing towards an unachievable state, and so expect them not to feel surprised and threatened by sceptical arguments.

The full-blooded sceptic could respond by pointing out that although many ideals do gain their functionality through the utility of being aspired towards, this is not required. To be functional, ideals need only promote progress towards some desirable state, even if that state is not the ideal itself – and even if no progress can in fact be made towards the ideal itself. A religious fictionalist might hold that God is an ideal that we orient ourselves towards as a means of progressing towards the ideals of love, justice or cohesive community (see e.g. Eshleman 2005). She might think this despite also thinking that since God does not actually exist, the ideal of God cannot itself be progressed towards. Similarly, the ideal of knowledge might promote progress along other epistemic dimensions, for example coherence or understanding, even if progression along these scales is not a progression along the scale towards knowledge.

The 'high standards' ideal sceptic also has replies available: scepticism might come as a threatening revelation not because it shows that we lack knowledge, but because it shines a light on how wide the gap is between the ideal epistemic state of knowledge and our current epistemic state. Analogously, say that justice is a regulative ideal which guides action. Although agents accept that their actions fall short of being maximally just, they take their actions to fall short in particular, identifiable ways. For example, someone might admit that their actions fail to be just on account of unconscious biases which lead them to favour some groups of people over others, while taking themselves to perform actions that are just insofar as they do not favour the agent over others in her community. Now, if this person were to be confronted with evidence that her actions typically fail to approximate justice along this second dimension, the dimension that she took herself to be doing well along, then this would be an alarming revelation. It would not be comforting to be reminded that since she never thought that her actions met the ideal of justice, she is no worse off, justice-wise. In the same way, someone who orients her epistemic practice towards the ideal of knowledge may take herself to have eliminated most error possibilities, while acknowledging

that biases or misleading evidence may have led her astray. Sceptical arguments which make salient the wide swathes of uneliminated possibilities threaten the progress that she took herself to have made towards the ideal. This is the case even if she did not expect to achieve the ideal.

Ultimately I think that there are three reasons to think that knowledge appears to fit better with typical examples of metaphors than with typical examples of regulative ideals. First, although Butchvarov has an expansive conception of regulative ideals which includes more concepts than is typical, on most views regulative ideals tend to denote concepts that are not employed in everyday contexts. Rather, regulative ideals tend to denote concepts that guide entire activities and domains, rather than individual actions and decisions within those domains. In the case of knowledge, knowledge as a regulative ideal might denote the end point of a large scale enquiry, for example, the quest of 'science' when seen as one, unified project. Thinking about each individual advancement in terms of approaching an ideal might frame these epistemic activities when seen at a remove. But thinking of each inquiry, however mundane, in terms of the ideal is perhaps not a typical use of ideals. Metaphor however is pervasive in shaping our thinking, both mundane and large scale.

Second, when employing ideals, it is plausible that agents typically take themselves to be aiming for the ideal, even if they fail to reach it, and even if reaching it is not required by the particular activity under consideration. Someone aiming to make a just law might very well use the regulative ideal of justice to make the most just law that she can. Yet on the ideals view that I have sketched, agents are not typically aiming for the ideal, at least not each and every time that the ideal is invoked. Rather, the scale towards the regulative ideal is the measure by which the assertability of "S knows that P" is adjudicated. Indeed, aiming for the ideal would be actively harmful, since it would waste resources. So in the case of knowledge, either knowledge is distinct from typical ideals, or agents are aiming for the ideal of knowledge, yet fail to realise this, which raises phenomenological concerns. In contrast, metaphor explains how a source domain shapes a target domain without positing that agents engaged in aiming at the target domain are also aiming at the source domain.

Finally, and more speculatively, one might worry that the regulative ideal of knowledge paves over the binary that makes knowledge functional in the first place. The ideals view emphasises the binary between the ideal of knowledge, and falling short of the ideal of knowledge. Yet the binary that really matters for making knowledge functional is the binary between being just insufficient for present purposes, and being barely sufficient for present purposes. On the view I have sketched, saying that S has achieved the ideal indicates that S is close enough to the ideal, which conveys that they have crossed the 'barely sufficient' line. But the YARDSTICK function is only realised if agents still at least to an extent keep the distinction between reaching the ideal and not reaching the ideal in mind. Yet it might turn out that the function typically played by knowledge does not require the YARDSTICK function - we do not typically need to reflect on the extent to which our epistemic position might be improved, and in what ways. Indeed, in many cases the time and energy spent on reflecting on ways to improve might come at a cost, even if ultimately these ways are dismissed as not needed. The metaphor view on the other hand frames the 'good enough' epistemic position in terms of knowledge, but need not require that agents consciously contrast their state to knowledge.

To summarise the chapter thus far, epistemic fictionalism understood in terms of metaphor fares better than three other non-face-value views: hyperbole, loose talk and regulative ideal accounts of knowledge talk. In particular, I have argued that seeing sentences of the form “S knows that P” as fictional better coheres with the function of knowledge talk, and better reflects our linguistic practices. I now turn to epistemic contextualism and sceptical pragmatic invariantism.

6.3 Epistemic Contextualism

Epistemic contextualism aims to explain the data which I used to motivate epistemic fictionalism - intuitions surrounding sceptical paradoxes, concessive knowledge attributions, and closure - but without positing a non-face-value mechanism. As such, epistemic contextualism might appear preferable to fictionalism.

According to epistemic contextualists, knowledge behave analogously to how face-value context sensitive terms such as 'here', 'always', and 'tall' behave. On epistemic contextualism, the semantic content of sentences of the form “S knows that P” varies according to context. In some contexts, sentences of the form “S knows that P” semantically express that S stands in a given knowledge relation with regards to P, and in other contexts they semantically express that S stands in a different knowledge relation with regards to P. Crucially, epistemic contextualist theories posit that the different propositions communicated by sentences of the form “S knows that P” are all face-value readings. That is to say, there are many face-value readings of sentences of the form “S knows that P”, and which one gets attributed to speakers is (in part) a feature of the context. The precise features of the context which bear on the assertability of knowledge attributions might include the attributor’s or subject’s stakes, worries, interests, saliences and attention, among other factors. As such, epistemic contextualism might appear to deliver similar results to epistemic fictionalism, without positing a face-value and a non-face-value reading. If epistemic contextualists are right and knowledge does indeed behave like face-value context sensitive terms, then this would present a challenge to the view that knowledge talk is typically metaphorical.

Furthermore, epistemic contextualism appears to share the advantages that I flagged for metaphor theory. It offers a simple resolution to sceptical paradoxes: the sentence “I know that I have hands” is true when uttered in the course of everyday life but false when uttered after a sceptical error possibility has been introduced. For epistemic contextualists this poses no contradiction: ‘knows’ picks out a more demanding epistemic standard in a context where a sceptical possibility has been raised than it does in an everyday context. Crucially, both of these readings of ‘knows’ are face-value readings. Admittedly, epistemic contextualists do commit speakers and hearers to widespread error in that speakers and hearers do not notice that ‘knows’ is a context-sensitive term. But just as in the case of metaphor, this error is neither uncharitable nor at odds with the phenomenology of using knowledge discourse, since many context sensitive terms are not consciously recognised as context sensitive (Blome-Tillmann 2008)⁸⁰. Epistemic contextualism also successfully answers the question of how knowledge attributions fulfil a function, and indeed some contextualists have defended epistemic contextualism on function-

⁸⁰ Although, see Schiffer (1996), Feldman (1999), Pritchard (2002), Hawthorne (2004), MacFarlane (2005), Stanley (2004), and (Bach 2005) for objections on contextualism that turn on the implausibility of the contextualist’s posited error.

first grounds⁸¹. In a nutshell, attributing knowledge serves to indicate that one is in a good enough epistemic position that members of one's epistemic community may safely draw on their knowledge, or such that one may draw on it for the specific purposes at hand, where these are higher. Whether one is in a good enough epistemic position depends on the context, and therefore, knowledge is context sensitive.

6.3.1 Linguistic Analogues

I want to start by considering the type of context-sensitive term that knowledge is supposed to resemble on the contextualist's picture. The once popular view that knowledge behaves like a gradable adjective (e.g. DeRose 1995, Cohen 1999) has been the target of sustained objections, and is now less widely endorsed. On this view, knowledge is analogous to gradable adjectives like 'tall' and 'empty'. Whether utterances of the sentences "Matilda is tall" and "the fridge are empty" are true depends in part on features of the context.

Jason Stanley (2005) has mounted a series of objections against the view that knowledge is a gradable adjective. Here are four ways in which knowledge does not appear to behave like a gradable adjective⁸².

Gradable adjectives admit of modifiers such as 'very' and 'quite', while knowledge does not. Compare:

- (1) x is very / quite / extremely / moderately flat.
- (2) S very / quite / extremely / moderately knows that P.

Gradable adjectives can be embedded in comparative and superlative constructions. Compare:

- (3) x is flatter / more flat than y
- (4) S knows that P more than y
- (5) x is the flattest F
- (6) S knows P the most.

Gradable adjectives admit of clarification as to which standards are operative in the context, while knowledge does not (see also Hawthorne 2004). Compare:

- (7) S is tall for a youth team player, but not tall for a footballer
- (8) S knows for a layman, but not for a scientist

Finally, gradable adjectives admit of context shifts mid-sentence, while knowledge does not. Compare:

- (9) The mouse is tall, but the elephant is not tall
- (10) S knows she has hands, but S does not know that S is not a handless brain in a vat

⁸¹ See Henderson (2011), Hannon (2013) and McKenna (2013)

⁸² See also Hawthorne (2004) and Davies (2005).

Some epistemic contextualists have responded that these objections miss the mark because they rest on a misunderstanding of epistemic contextualists' key commitment. The objections target the claims that knowledge is not a gradable adjective and does not come in degrees, but contextualists need not endorse these claims. Rather, the epistemic contextualist thinks that knowledge behaves analogously to gradable adjectives *in certain respects*. Blome-Tillman (2008) and McKenna (2015), for example, argue that the relevant similarity between knowledge and gradable adjectives is that knowledge is semantically linked to an epistemic scale, in particular that knowledge has a "semantic link to a scale of epistemic strength" (Blome-Tillman 2008:31). They argue that knowledge is far from unique in this respect. *Snore* and *shout* are other examples of concepts that are linked to scales, for example, to loudness scales. The sense in which knowledge is gradable is the sense in which 'snore' is gradable: instances of snoring and knowledge can be ordered along a scale according to how loud, or how strong, they are.

Blome-Tillman argues that making explicit the scale that knowledge is semantically linked to avoids the first two objections raised by Stanley. The following sound just fine:

- (11) S snores very / quite / extremely loudly
- (12) S knows that P with very / quite / extremely good evidence
- (13) S snored more loudly than R
- (14) S knows that P with better justification than R knows that q

Likewise, making the scale explicit admits of clarification, for we are now in a position to clarify the scale-relative standards regarding which the occurrence in question is being evaluated:

- (15) S is shouting loudly for someone at the library, but is not shouting loudly for someone at the pub
- (16) S knows that P on the basis of evidence good enough for a layman, but not on the basis of evidence good enough for a scientist

Mid-sentence context shifts are admittedly still not permitted on this picture, but Blome-Tillman handles this objection by denying that easily recognisable mid-sentence shift is a universal characteristic of context sensitive terms. The following dialogue reveals contextual shift blindness in an uncontroversially context sensitive term:

- A: That warehouse is empty.
- B: But have you considered there are some dust particles in there?
- A: I guess I haven't.
- B: So you admit that warehouse isn't empty, and so you were wrong earlier?
- A: Oh, c'mon! I didn't say that the warehouse is empty.

A's response seems out of place. Yet on an epistemic contextualist reading, every sentence in the dialogue has a true reading. The appearance of contradiction is a result of speakers and hearers being unaware of a shift in the extension picked out by the context sensitive term 'empty'. Yet,

such unawareness does not show that ‘empty’ has invariant semantics. More plausibly, it shows that whether a term is easily identifiable as context sensitive is not a necessary condition of it being context sensitive.

There is however one problem: it is not obvious that ‘snore’ and ‘shout’ are best understood as context-sensitive terms. It is compelling that these terms are semantically linked to a scale (e.g. of loudness) in the sense that instances of snoring can be ordered along the scale, and that if they fall sufficiently low down on the scale, they cease to count as instances of snoring. And undeniably, one can snore more or less loudly. However, this is not to say that a point in the scale which does count as snoring in one context might not count as snoring in a different context. After all, the moderate invariantist also thinks that people snore and shout more or less loudly, and that people know on the basis of better or worse evidence. This is the distinctive claim that the contextualist needs to establish: not just that there is a semantic link to a scale, but that whether a particular occurrence counts as an extension of the term is in part dependent on where that instance falls on the scale that the term is semantically linked to.

McKenna (2015) tries to motivate the intuition that snore and shout are context-sensitive by spelling out the details of sentence (15). Say Kevin is engaged in philosophical discussion in the library and raises his voice beyond the levels acceptable in the library. A student walks by and says “That guy is shouting. I wish he would stop!”, but the same student could overhear the same conversation held at the same volume in the pub and say truly “He isn't shouting, he's just talking to his friend”. Similarly, if Kevin falls asleep in the library and starts breathing more heavily than normal the student might truly say “He's snoring. I wish he would stop!”. Yet when Kevin is asleep in bed making the exact same sound, his wife might say “He's not snoring, he's just asleep”. The intuition is supposed to be that both the student and Kevin's wife say something true. Not only do features of the context partly determine whether a particular sound falls is properly described as ‘snoring loudly’ in that context, but so too do features of the context partly determine whether it is properly described as ‘snoring’ in that context.

The worry with this argument is that so spelled out, these examples are best interpreted as instances of hyperbole than context sensitivity. First, the student is not aiming to simply describe the situation. Rather, she is using that particular term to mark her displeasure at the inappropriate loudness of the sound that Kevin is making. As discussed in §6.1, using maximal standards to convey one's attitude to a situation is precisely the function that hyperbole plays. Second, it seems that it would be very easy to force the student in the library to retract her claim. Indeed, even the gentlest of challenges would do it – “he isn't snoring, he's just breathing heavily” or “Yes he's talking way too loudly - but he's not *shouting*, is he?”. These challenges do not alter the context in particular ways. They do not raise the stakes, or force the student to pay attention to previously unnoticed features of the context, and the same challenges would not produce retraction in the pub scenario. More plausibly, snoring and shouting are not context sensitive. They are invariant terms with indeterminate cases at the edges, and which can be used hyperbolically to convey that a sound is too close to snoring and shouting to be acceptable in the present context.

A related strategy which avoids this problem is to simply accept that knowledge is analogous to ‘snoring loudly’ rather than ‘snoring’. That is to say that knowledge is analogous to modified

gradable adjectives, like 'tall enough', 'reasonably empty', 'very loud'. This view is defended by Halliday (2007) and Jaster (2013). It retains the gradable adjective features of Blome-Tillman's account, such as the emphasis on the semantic link between context sensitive terms and a scale, but it adds that knowledge functions as a pure threshold predicate: its sole purpose is to mark "whether the threshold is surpassed or not" (Jaster 2013:322). The expression 'tall enough' marks the threshold at which one is sufficiently tall for present purposes, and is linked to a scale: it is possible to order, by tallness, all of the people who satisfy the condition of being tall enough. Yet being tall enough does not come in degrees. One either is, or is not, tall enough.

Modified gradable adjectives are unproblematically gradable in the sense of being linked to a scale, yet they also fail three of Stanley's gradeability tests. Therefore, that knowledge also fails Stanley's tests is no objection to this version of contextualism. The following sentences are just as unpromising as the knowledge analogues:

- (17) S is very / quite / extremely / moderately tall enough
- (18) S is more tall enough than R is tall enough
- (19) S is the most tall enough
- (20) The mouse is tall enough, but the elephant isn't tall enough

Admittedly, if knowledge is a modified gradable adjective we might expect it to pass the clarification test, for the following is perhaps a paradigmatic use of 'tall enough':

- (21) Marta is tall enough for the spinning teacup ride, but not tall enough for the rollercoaster

However, some epistemic contextualists hold that knowledge does pass this test, once the scale to which knowledge is semantically linked is spelled out (Ludlow 2005:20). Blome-Tillman (2008) argues that knowledge fails this test on account of being a gradable verb rather than a gradable adjective.

Stanley objects that modified gradable adjectives are bad analogues for knowledge because "one standard use of these expressions is to convey that something has the property for a sufficient degree for present purposes, though it does not in fact have the property." (p. 44). Stanley worries that "S is tall enough" typically implies that S is not actually tall, although they approach being tall to a sufficient degree that their lack of tallness is not an obstacle to them performing the present task. This is a problem because clearly a standard use of 'knows' is not to convey that one is not justified enough. Further, the following sentence is non-standard, and hard to make sense of:

- (22) She isn't in a strong enough epistemic position with regards to P, but she does know that P.

However, as Jester points out, Stanley's objection ignores that knowledge only applies to relatively strong epistemic positions: even contextualists can say that knowledge only applies to cases that have crossed a sufficiently high threshold. In this regard, knowledge is like 'very tall', 'extremely

flat', and 'excellent'. The sentence "S is very tall, but isn't tall" or "S is excellent, but isn't good" are as hard to make sense of as (22). The right analogue, then, is:

(23) "S isn't in a maximally strong epistemic position, but she does know"

I think that as far as linguistic analogues are concerned, modified gradable adjectives are the best model for context-sensitive semantics of 'knows'. But we might wonder at this point whether the contextualist explanation is indeed more straightforward or in possession of more theoretical virtues than the metaphor theory. For knowledge is clearly not a modified gradable adjective. As Jester rightly points out, objections to epistemic contextualism that aim to show that knowledge is not an example of that exact type of linguistic type completely miss the mark: the point is rather that knowledge is analogous to this type, in virtue of sharing the key context sensitive characteristic of this type. In contrast, on the metaphor view, knowledge *is* a metaphor. It functions just like metaphors, because it is one. This is simpler than saying that it uniquely among verbs behaves like a modified gradable adjective. Why would knowledge behave like a modified gradable adjective, without being a modified gradable adjective? This is not to say that such a view is logically incoherent. It could be that knowledge is unique in this sense. However, we should not posit uniqueness while there are other options on the table. As McKenna (2015:365) puts it, "if the contextualist is forced to posit *sui generis* context-sensitivity, she's in real trouble, especially because there are other views that can account for the linguistic usage that provides some support for contextualism".

The view that knowledge is a verb with a semantic link to a scale like snore and shout would have been a threatening alternative if it could be shown that these verbs actually are context sensitive. But if the best context-sensitive model is an analogue rather than the category that knowledge properly belongs to, metaphor theory has the linguistic advantage in being able to categorise knowledge.

6.3.2. Everyday Context Shifts

Still, the epistemic contextualist might reply that understanding knowledge as a context-sensitive term has one crucial advantage: it explains variability in everyday knowledge attributions. So far I have focused on context-sensitivity as generated by sceptical paradoxes, arguing that sentences of the form "S knows that P" are false if taken at face-value but nevertheless succeed in communicating a true, non-face-value content. But the more typical cases of context-shift involve raising the standards to ones that are more demanding than the present standards, while still falling dramatically short of maximal standards.

Consider this standard example (Stanley 2005:4, adapted from DeRose 1992):

BANK

LOW. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it

isn't very important that their paychecks are deposited right away, Hannah says, "I know the bank will be open tomorrow, since I was there just 2 weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning".

HIGH. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank 2 weeks before on Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, "I guess you're right. I don't know that the bank will be open tomorrow".

The epistemic contextualist diagnosis of why we want to say that only **Low-Hannah** knows is that **High-Hannah** finds herself in a context where the standards for true knowledge ascriptions are higher than those operative in **Low-Hannah's** context. There is no contradiction between **Low-Hannah's** knowledge attribution and **High-Hannah's** denial of knowledge because the term 'knows' picks out different standards in these different contexts.

What about epistemic fictionalism? The diagnosis of context-sensitivity used to explain attitudes to sceptical arguments does not neatly apply to these cases. The fictionalist can say that **Low-Hannah's** knowledge attribution is true if given a non-face-value reading and false on a face-value reading. But she cannot say that **High-Hannah's** initial knowledge attribution is false in virtue of being false at face-value. It *is* false at face-value, of course. But unlike **Low-Hannah's** knowledge attribution, **High-Hannah's** knowledge attribution is also false on a non-face-value reading. The distinction between face-value and non-face-value readings is insufficient for accounting for explaining everyday context variability.

Chung remarks that it is notoriously difficult to spell out the content of non-literal utterances, but that "since non-literal speech is often deeply context sensitive, we should expect the answer to this question to be more or less context-sensitive as well, depending on the pragmatic machinery involved" (2018:391). Metaphor is particularly context sensitive, and I think that it does have the resources to explain everyday context shifts. I have argued that metaphor frames structure the target concept. However, whether a metaphor is apt depends on features of the context. Additionally, which aspects of the source domain structure the target domain is sensitive to features of the context (Camp 2006, Wearing 2008).

The **BANK** case illustrates how a challenge from an interlocutor can raise the standards for knowledge attributions by changing the context to one where higher standards are operant. In **BANK**, this challenge takes two forms: the stakes are raised, and the hearer challenges the speaker's knowledge assertion by raising an error possibility.

Imagine this conversation between Romeo and his cousin Benvolio:

UNIQUE

R: Juliet is the sun

B: She's a Capulet. Be smart about this: is she really the sun, is there really no other woman in all of Verona, maybe one whose cousin isn't always trying to kill you?

R: Ok fine, I guess you're right, she's not the sun

In chapter 4, I discussed the case in which Benvolio raises the standards operative in the context by challenging the face-value content communicated by Romeo's utterance, thus forcing Romeo to own or retract the face-value reading. But in UNIQUE, Romeo is not ceding to Benvolio on the grounds that Benvolio is forcing Romeo into a context where he must answer to the face-value implications of his utterance. Rather, Benvolio is challenging the *aptness* of Romeo's metaphor. That is to say, Benvolio is challenging the extent to which the source domain (the sun) does in fact frame the target domain (Juliet). Benvolio does this by challenging the extent to which Juliet is in fact unique. So, he challenges whether one particular salient dimension of the source domain (the uniqueness of the sun) applies to the target domain. In drawing Romeo's attention to the animosity that Juliet's cousin has towards Romeo, Benvolio is also making salient that the stakes are high.

On many standard formulation of metaphor, the source domain frames the target domain in part by emphasising salient similarities between the source domain and the target domain, which engenders an altered perception of the target domain⁸³. These similarities might be properties, relations, structures or some other more indirect feature (Camp 2009). Similarity is uncontroversially a context-sensitive notion. Whether A is similar to B depends on features of the context, including the dimension along which A is similar to B and the degree of similarity. For example, "honeybees are similar to bumblebees" might be true according to the standards and purposes operative in a primary school classroom, but false according to the standards and purposes operative in a Scottish wildlife preservation programme. It is not just that there are dimension of similarity that matter in the preservation programme but not the classroom (e.g. differences in habitat and honey production) but that the degree to which honeybees and bumbles are similar along the same dimension might be sufficient to class them as similar in one context but not the other.

If the aptness of a metaphor depends on whether the source and target domain are sufficiently similar in the key salient respect, then metaphor aptness is straightforwardly context sensitive. Whether a metaphor is apt will depend on whether the standards operative in the context make it the case that the source and target domain are sufficiently similar. In the conversation above, Benvolio succeeds in raising the standards operative in the conversation by pushing a dissimilarity between the source domain (the sun) and the target domain (Juliet), namely, that the sun is unique on account of there being only one sun, while Juliet is not, on account of there being other ladies in Verona. Furthermore, Benvolio raises the stakes by highlighting the danger that

⁸³ See Miller (1979) Searle (1969), Davidson (1978), Camp (2009) On some views, the correct face-value paraphrase for a metaphor of the form 'S is R' is 'S is similar to R, in regards to x'. On category views, the metaphor creates a new ad hoc category which has as extensions both the extensions of the source and target domains. Although this view does not explicitly involve similarity, metaphor creates a frame by re-categorising the target domain into a new category. The extent to which the category is coherent will also depend on the extent to whether there is a salient similarity between the source and target domains.

Juliet's cousin poses. Whether or not the similarity is sufficient to survive the raising of the potential dissimilarity depends in part on features of the contexts, in particular, on the purposes and stakes.

On the metaphor picture that I have defended, the source domain of infallibilism frames our current epistemic position. Whether "S know that P" is an apt metaphor for framing the epistemic position that S is in with regards to P is in part a function of how similar S's epistemic position with regards to P is to infallibilist – and this similarity in turn is a function of the context. In BANK, Sarah pushes a dissimilarity between the source domain (infallible knowledge) and the target domain (Hannah's epistemic position), while raising the stakes. As a result, Hannah's metaphorical assertion "I know that the bank will be open tomorrow" is no longer apt.

An interesting feature of metaphor is that it allows for a second type of context-sensitivity: whether a metaphor is apt depends in part on whether the dimension that is salient in the context is apt. Camp (2009) argues that Romeo's claim that "Juliet is the sun" communicates a variety of related propositions, and features of the context emphasise which ones are particularly salient in that context. In UNIQUE, Benvolio challenges the extent to which 'Juliet is the sun' is apt, by challenging whether Juliet is unique. But he could also challenge Romeo's claim by putting pressure on, and thus making salient, other dimensions of the mapping from source domain to target domain. Consider these modifications.

CENTRE OF UNIVERSE

R: Juliet is the sun

B: She's a Capulet, be smart about this: Is she really the sun, you've only known her five minutes, and her cousin is always trying to kill you.

Here Benvolio challenges the extent to which Juliet is the centre of Romeo's life, by making salient that he has only known her a very short amount of time. He is challenging the aptness of the 'centre of the universe' dimension of the sun for framing Romeo's relationship to Juliet as the principal preoccupation of his thoughts by implying that it takes time for a person to become the centre of another person's life.

GROW

R: Juliet is the sun

B: She's a Capulet, be smart about this: Is she really the sun, she's like what, 13 years old? And her cousin is always trying to kill you.

Here Benvolio challenges the extent to which Juliet is well positioned to help Romeo grow and mature by making salient that Juliet is herself very young. He is challenging the aptness of the 'growing and nurturing' dimension of the sun as an apt frame for Juliet.

I would like to tentatively suggest that this type of context-sensitivity is also present in how knowledge attributions are challenged. In BANK, Sarah challenges Hannah along what I have

argued is a typically salient dimension for knowledge frames: Sarah challenges whether Hannah is in a strong enough epistemic position to act as an informant with regards to whether the bank is open. However, there are other features of knowledge that we might expect to go hand in hand with strong epistemic positions such as psychological certainty and availability of strong evidence. Furthermore, as Craig (1990) argues, there are good reasons to expect these features to go together. Agents who are not psychologically certain with regards to P or who cannot produce strong evidence in favour of their belief that P will not typically be good enough informants that P. So, we can expect sentences of the form “S knows that P” to typically imply (if not outright communicate) these related propositions, that S is psychologically certain that P and that S has good evidence that P. We can imagine challenges to Hannah's knowledge attribution in the high standards context that put pressure on these dimensions, analogously to how Benvolio's various challenges put pressure on different dimensions of the aptness of "Juliet is the sun". Consider:

CERTAINTY

H: I know that the bank is open on Saturday

S: Do you feel absolutely certain?

H: Well now you've said that I'm doubting myself

EVIDENCE

H: I know that the bank is open on Saturday

S: How do you know - did you check the hours online this morning?

H: Well no, but I was here two weeks ago.

In CERTAINTY, Sarah challenges the non-face-value truth of Hannah's knowledge attribution. Specifically, she challenges the aptness of the frame between source and target domain by challenging the extent to which the dimension of being maximally psychologically certain does in fact apply to the target domain. In EVIDENCE, Sarah challenges the aptness of the frame between source and target domain by challenging the extent to which the dimension of being supported by good evidence does in fact apply to the target domain. If it is right that metaphor is context-sensitive in this second, dimension-related way, then metaphor theory has a neat account of why these different and seemingly unrelated challenges might result in context variability across contexts.

6.3.3. A complication: is metaphor underpinned by contextualist mechanisms?

I have argued that a metaphor based epistemic fictionalism fares better than epistemic contextualism when it comes to interpreting the linguistic data, and just as well as epistemic contextualism on explaining everyday context variability. There is a complication to the metaphor picture, however, concerning the ongoing debate in metaphor theory about the mechanism which underpins metaphor. In particular, the debate concerns how exactly uttering a metaphorical sentence succeeds in communicating the relevant non-face-value content. Two popular

responses to this question are that metaphor operates on implicature and that metaphor is context-sensitive⁸⁴. On the implicature view, hearers first parse metaphorical sentences at face-value, and only when the face-value content violates a conversational maxim (truth or relevance, typically) access the implicated metaphorical communicated content. On this view, the face-value content of a metaphorical utterance is its only semantic content, and the content that is asserted by utterances of metaphorical sentences. The non-face-value content is communicated by conversational implicature. In contrast, on the contextualist view, metaphorical sentences enjoy more than one semantic content: the face-value content, the non-face-value content, and perhaps others. Which content gets activated is determined in part by features of the context. On this view, any given non-face-value content is one of at least two, and typically many, semantic contents, and is directly asserted by speakers.

The complication, then, is that positing that knowledge is metaphorical might appear to kick the real question down the road - that is, the question of explaining context variability in what sentences of the form “S knows that P” communicate. I’ve said that metaphor gives us a nice way of explaining this variability, but as it turns out, the question of how metaphor does this is left open - and worse, the principal options in this debate are precisely those familiar from the debate between epistemic contextualists and their critics. If we favour the contextualist view of metaphor, we might wonder what explanatory powers are gained by positing metaphor to begin with, as opposed to being traditional face-value epistemic contextualists. On a metaphor contextualist picture of knowledge, to say “S knows that P” semantically expresses at least two contents, the face-value one and the non-face-value one. Which one is active in any one context depends in part on features of the context. But the epistemic contextualist might well object that epistemic contextualism already generates this result, minus the appeal to metaphor and the distinction between face-value and non-face-value content. On the other hand, the arguments that favour implicature views in metaphor theory are the very ones that epistemic contextualists seek to resist when applied to knowledge, and so those persuaded by epistemic contextualism are unlikely to favour an implicature view of metaphor.

There are two responses to this concern, the first highly conciliatory: that it is open to the contextualist to take metaphor as one possible candidate underlying semantic mechanism, alongside ‘gradable adjectives’ and ‘indexical’. Indeed, a contextualist troubled by apparent dissimilarities between knowledge and these other context sensitive categories might welcome another option, especially one that is already as entrenched in everyday language as metaphor is. On this view, epistemic fictionalism would be a version of epistemic contextualism.

Indeed, appeal to metaphor theory explains why knowledge fares poorly on the tests for context sensitive terms discussed in §6.3.1.. On this view, metaphorical sentences are uncontroversially context sensitive, and yet do not admit of modifiers. The following for example do not work:

- (24) Juliet is very / quite / extremely / moderately the sun
- (25) My lawyer is very / quite / extremely / moderately a shark

⁸⁴ See Grice (1989), Martinich (1984), Camp (2006) Arsenault (2020) for implicature accounts of metaphor. See Crimmins (1998), Hills (1997), Carston (2002), Recanati (2004), Bezuidenhout (2001) for contextualist accounts of metaphor.

(26) That surgeon is very / quite / extremely / moderately a butcher.

Yet metaphor is context sensitive, and allows modifiers to specified dimensions of similarity:

(27) Juliet is the sun. She is extremely beautiful.

(28) My lawyer is a shark. She's really ruthless.

(29) No, my surgeon isn't a butcher - he wasn't at all careless

The same is the case for knowledge talk. The following read fine:

(30) S knows that P. S's evidence for P is extremely good.

(31) S doesn't know that P. S's evidence for P isn't at all good.

Neither do metaphors readily admit of comparatives and superlatives. The following does not work:

(32) Juliet is more the sun than the other ladies of Verona

(33) My lawyer is the most shark

But again, we can easily apply comparatives and superlatives to specified dimensions of similarity between the target and source domains:

(34) Juliet is the sun, and the other ladies of Verona are the moon. Juliet is more radiant than the other ladies of Verona.

(35) My lawyer is a shark. She's the most ruthless person I've ever met.

Metaphor does not lend itself to clarification regarding the standards that are in operation. We cannot say:

(36) Juliet is the sun for a girl, but not the sun for a celestial body

(37) My lawyer is a shark for a person, but not a shark for an animal

But when speaking metaphorically, one can clarify whether they intend to communicate the face-value or non-face-value content of their utterance:

(38) Oh, I don't mean Juliet is literally the sun. I meant that she's radiant.

(39) No, I'm not talking metaphorically, the lawyer in that cartoon is an actual shark

One can also clarify the salient dimension along which the metaphor is apt:

(40) Bill is a bulldozer. Not in the sense that he's pushy, but he never admits defeat in the face of an obstacle

Likewise, although such sentences are not found in everyday speech, in knowledge talk we can in principle clarify whether we intend to communicate the face-value or non-face-value content of an utterance:

(41) S knows that P. Not in the sense that S can rule out the possibility that S is a BIV, but in the sense that S's epistemic position with regards to P is good enough

And we can also clarify the salient dimension along which the metaphor is apt:

(42) I know that the bank is open. Now that we're talking about it I'm starting to feel unsure, but I do clearly remember it being open two weeks ago.

Finally, as I argued in chapter 4, metaphor does not admit of mid-sentence shift. For once the face-value reading of a fragment of discourse is in play in a conversational context, then this reading carries over to the entire sentence. It seems, then, that metaphorical sentences also fail to pass the tests for context-sensitivity. Yet this does not indicate that metaphorical sentences have only one meaning. Rather, it shows that the context-sensitivity of metaphorical sentences does not come out linguistically in the same way as non-metaphorical sentences. Therefore, the epistemic contextualist who is disturbed by the difficulty of fitting knowledge talk into a category of uncontroversially context-sensitive terms might welcome the suggestion that knowledge talk is context-sensitive because metaphorical.

The second, less conciliatory response that the epistemic fictionalist can give to the epistemic contextualist is that epistemic fictionalism is not explanatorily inert, even on a contextualist view of metaphor theory. *Contra* what some critics of epistemic contextualism have argued (e.g. Pritchard 2018, Wright 2005), contextualism is not a theory which privileges any one epistemic standard over another. According to epistemic contextualism, whether a sentence of the form "S knows that P" denotes one epistemic standard over another in any particular context is a function of the standards at play in that context. Crucially, there is no such thing as a context-free reading of "S knows that P". No standard is "more proper" (Davis 2015:36) than any other, but rather each standard is proper in a context only insofar as it is the standard operative in that context. This means that on epistemic contextualism, there is nothing more correct about the sceptic's standards. Indeed, in most contexts, the sceptic's standards are improper.

However, this means that epistemic contextualists must appeal to an error theory to explain why the standards of a context are so easily raised so that one can, with ease, find themselves in context in which "S knows that P" expresses a false proposition, but conversely only with difficulty bring the standards back down. In particular, they need to posit an error theory to explain why the sceptic's standards are very easy to induce, but only with difficulty dismissed again. If all contexts are equal from the perspective of no-context, why does the sceptic so easily succeed in inducing her context? On the flip side, why does the anti-skeptic struggle to re-induce the context in which sceptical standards do not apply?

On epistemic fictionalism, "S knows that P" at face-value communicates what the sceptic demands, and so the context where most utterances of this sentence are false can easily be induced. In contrast, once the face-value reading is operative in a context, it is very difficult to reinstate the non-face-value reading when in conversation with a non-cooperative interlocutor. In general, more demanding standards are closer to the face-value reading and share more similarities with the source domain. It is thus far easier to introduce into a conversational context

possibilities which push a dissimilarity between the target domain and the source domain, than it is to introduce a similarity. Epistemic fictionalism therefore explains why it is easier to raise the demandingness of standards operative in a context than it is to lower them.

6.4 Sceptical Pragmatic Invariantism

Like epistemic contextualism, sceptical pragmatic invariantism is another view which explains context variability by appeal to a widespread and relatively uncontroversial linguistic mechanism, and as such might at first blush appear preferable to epistemic fictionalism. This section aims to motivate the claim that those sympathetic to sceptical pragmatic invariantism would do well to endorse epistemic fictionalism.

According to sceptical pragmatic invariantism (henceforth: *SPI*), utterances of sentences of the form “S knows that P” assert a proposition that is typically false, for example, that ‘S is infallible with regards to P’. Nevertheless, many such sentences are assertable because despite being false, they *pragmatically implicate* something true, perhaps, that S is justified in believing that P, or has ruled out all of the relevant alternatives to P. According to *SPI*, the content implicated by an utterance is in part fixed by the context, and therefore, whether a sentence is assertable varies with context. Here is how Dinges (2016:2579) characterises *SPI*’s explanation of context variability in cases like **BANK**, using the character of DeRose as the protagonist:

DeRose indeed fails to know that the bank will be open in the bank cases. This means that his knowledge denial in the high standards case is literally true and that his knowledge ascription in the low standards case is literally false (because invariantism is correct). The knowledge denial seems appropriate (or even true) simply because it is true. The knowledge ascription seems appropriate (or even true), not because it is true, but because it has a true implicature. These are the basic ideas behind *SPI*.

Understanding context variability in terms of variability in implicated content is not unique to sceptical invariantism. Indeed, several non-sceptical invariantists also spell out variations in intuitions regarding assertability in terms of the general linguistic mechanism of Gricean implicature⁸⁵. According to Griceans, speakers communicate implicated content by getting hearers to recognise that the face-value reading of their utterance would violate a conversational maxim. The hearer reconciles the speaker’s utterance with the general expectation that the speaker is nevertheless co-operative, i.e., communicating something relevant and informative, if not at the level of what is strictly asserted, then at the level of what is implicated, extracting the proposition communicated by (one of) the non-face-value reading of the speaker’s utterance. The speaker wishes to communicate a specific content to the hearer via implicature and aims to get the hearer to work out what that content is, trusting that

⁸⁵ See for example Rysiew (2001, 2007), Dougherty and Rysiew (2009) Brown (2006, 2018), Black (2005), Bach (2005). The pragmatic sceptical invariantist holds that knowledge attributions are typically false, but in everyday contexts are nevertheless assertable on account of their implicating something true. In contrast, the moderate pragmatic sceptical invariantist holds that many knowledge attributions are true, but that nevertheless some of these true knowledge assertions are unassertable in particularly demanding contexts, on account of implicating something false.

the hearer will make the right inferences based on the utterance provided in the context. In Grice's (1975) classic example, a philosophy professor tasked with writing his student a letter of recommendation for a philosophy post writes that the student is punctual and has good handwriting. If taken at face-value, such a communication would violate the conversational maxim of relevance: the student's handwriting does not qualify them for a job in philosophy. The letter writer expects the reader to retrieve the implicated content that the student is a poor philosopher. In the same way, proponents of SPI posit that speakers assert the false proposition "S knows that P" in order to implicate some other, typically true content.

As discussed in chapter 5, a significant portion of metaphor theories classify metaphor in general as implicature. So just as was the case for epistemic contextualism and an epistemic fictionalism underpinned by contextualist theories of metaphor, not all versions of epistemic fictionalism are in conflict with sceptical pragmatic invariantism. For example, a force epistemic fictionalism based on Camp's implicature metaphor theory would hold that uttering sentences of the form "S knows that P" asserts a false face-value content – S is infallible with regards to P – in order to thereby implicate some other proposition, e.g. that S is a good informant with regards to P. This version of epistemic fictionalism would be a version of sceptical pragmatic invariantism, rather than a competitor to sceptical pragmatic invariantism⁸⁶.

Nevertheless, mirroring the epistemic contextualism objection, the sceptical pragmatic invariantist may ask: why think that the pragmatic mechanism is that of metaphor, as opposed to straightforward Gricean implicature? What does epistemic fictionalism add to the explanation of how sentences of the form "S knows that P" succeed in communicating their content? The remainder of this chapter argues that epistemic fictionalism is well placed to respond to two objections that straightforward non-metaphor Gricean sceptical practical invariantism faces. To illustrate this point, I compare standard Gricean SPI to an epistemic fictionalism on which the fictional mechanism is metaphor, on a Gricean implicature understanding of metaphor. I also comment briefly on how non-implicature based metaphor fictionalism might address these two objections, but the focus will be on a Gricean implicature metaphor fictionalism. This focus should not be read as endorsing Gricean implicature as the correct account of metaphor, or even of metaphor fictionalism. Rather, the aim is to show how if one is already persuaded by standard SPI, then spelling out the pragmatic mechanism in terms of metaphor confers advantages. That is to say, I hope to make plausible the claim that if one is already persuaded by standard SPI, then one has good reason to become a fictionalist.

⁸⁶ On many theories of loose talk, loose talk also communicates content via implicature. Indeed, Dinges writes that "proponents of SPI typically hold that the knowledge claim in the low standards scenario is a case of *loose use*" (2016:2580). On this view, uttering a sentence of the form "S knows that P" implicates that S is close enough for present purposes to knowing that P. Since loose talk is discussed in §1 of this chapter, I won't discuss it further in this section, but do wish to note that the discussion on implicature in general also applies to theories of loose talk on which loose talk conveys its communicated content via implicature.

6.4.1. Knowledge attributions in thought

Baumann (2011) argues against both moderate and sceptical pragmatic invariantism by pointing to the fact that sentences of the form “S knows that P” are equally as felicitous when they are believed or thought, without being asserted. Imagine that it is the low stakes case: Hannah has gone to the bank alone, and seeing the queue, she forms the belief that the bank will be open on Saturday, and from this she forms the belief that she will be able to deposit her cheque tomorrow. Hannah’s belief appears felicitous in the low stakes case. The sceptical invariantist analysis of the original case was that Hannah asserts something false (e.g. ‘my evidence entails that the bank will be open tomorrow’) in order to implicate something true (e.g. ‘my evidence is sufficiently strong such that we may rely on the bank being open tomorrow’). But this implicature-based analysis does not apply at the level of belief, since beliefs do not carry implicatures. Beliefs cannot flout conversational maxims as there is no audience to go searching for an alternative non-face-value meaning. So, sceptical pragmatic invariantism cannot explain why beliefs of the form “S knows that P” appear appropriate and felicitous in most contexts, despite being false⁸⁷.

Metaphor fictionalism is well-equipped to explain knowledge attributions in thought. Metaphor is a thoroughly cognitive phenomenon, framing thoughts and perceptions as well as language (Lakoff and Johnson 1980, Thibodeau and Boroditsy 2011, Semino et. al. 2018). And indeed, almost all metaphor theories posit that metaphorical content is expressed in thought as well as in language. Non-implicature theories of metaphor have a straightforward reply to the cases raised by Baumann: metaphorical meanings are lodged into the semantics of metaphorical sentences, and so metaphorical meanings are expressed in both thought and assertion. In the low standard BANK case, the belief “I know that the bank is open” metaphorically communicates the non-face-value content “my evidence is sufficiently strong such that we may rely on the bank being open tomorrow”, rather than the face-value content ‘my evidence entails that the bank will be open tomorrow’. As discussed in chapter 5, that metaphorical sentences might contain a tacit operator, or perhaps they contain context-generated ad hoc metaphorical concepts.

The case is not so straightforward for implicature theories of metaphor. Indeed, Baumann’s objection could be reformulated into an objection to implicature theories of metaphor in general: implicature metaphor theories of metaphor posit that in uttering a metaphorical sentence speakers say something false, to implicate something true. But this analysis does not explain how it is that metaphorical sentences in belief and thought succeed in communicating the non-face-value content. The implicature theorist says that in uttering the sentence “Juliet is the sun”, Romeo implicates that she is radiant and worthy of worship. But it seems that Romeo’s belief or thought that Juliet is the sun communicates the same content as his utterance of “Juliet is the

⁸⁷ Although see Dinges (2016) for an invariantism-friendly solution.

sun”⁸⁸. How is this possible, if detecting conversational maxim flouting is what enables the expression of the metaphor’s pragmatically implicated non-face-value content?⁸⁹

An initially plausible solution might be to think of thought as a conversation with oneself, governed by the same conversational maxims as a conversation. However, this suggestion is not in keeping with how implicature is supposed to work. Implicature works because speakers trust hearers to work out that a conversational maxim has been broken, and trust them to search for a likely alternative meaning. This does not apply to conversations with oneself – one would already be aware they are using speech to communicate its non-face-value content. Speaks (2008:113) goes as far as to write into the definition of an implicature that it precludes the sentence being used to communicate the implicated content in thought: “the fact that a sentence S may be used in conversation to communicate (convey, assert) P can be explained as a conversational implicature only if S cannot be used by an agent in thought to judge (think) P.”

A better response is for the metaphor theorist to argue that metaphor in thought is a matter of what using the metaphor enables the thinker to achieve, that they could not otherwise have achieved (or could not otherwise *as easily* have achieved) using exclusively face-value meanings. In general, we should not expect language in thought to function in precisely the same way as language in conversation, for the latter is always shaped by what the speaker is attempting to get the hearer to read into their utterances. Griceans in particular think that *all* intrapersonal spoken communication, not just implicature and metaphor, is governed by adherence to conversational maxims. The Gricean metaphor theorist posits that when using metaphors in speech, a speaker utters a metaphorical sentence such as “You’re the cream in my coffee” in order to communicate to the hearer something akin to “You make me feel proud and happy”. Metaphor in spoken language is a matter of a speaker getting the hearer to believe some proposition (e.g. ‘you make me feel proud and happy’) from some other (albeit, false at face-value) proposition communicated by their utterance (e.g. ‘you are the cream in my coffee’). Uttering the latter proposition succeeds in communicating the former proposition because the hearer recognised that the speaker’s utterance, if taken at face-value, would violate a conversational maxim of relevance or quality. This sparks engagement into an inferential process such that the hearer, too, comes to see how they (as the subject) can play a role to the speaker that is somehow structured along the lines of cream in coffee. Perhaps, the speaker gets the hearer to recognise some salient similarity between the role that the hearer (as the subject) plays in the speaker’s life and the role that cream plays in a cup of coffee, or perhaps the speaker calls to the hearer’s mind common associations with cream in coffee.

Analogously, the Gricean can say that metaphor use in private thought is a matter of the thinker coming to believe some proposition (e.g., ‘The subject is my pride and joy’) from reflecting on some other (albeit false at face-value) proposition (e.g., ‘The subject is the cream in my coffee’). As is the case in the interpersonal spoken context, the thinker exploits the false proposition

⁸⁸ As it happens, this is close to Romeo’s situation. When he utters this sentence, he is alone and (excluding the play’s audience) has no interlocutors to detect any maxim flouting and retrieve the implicated content.

⁸⁹ Nogales (2012: 999) discusses a related point, noting that “we are perfectly able to interpret a metaphor found on a scrap of paper, despite a complete lack of knowledge of illocutionary features, something that would be impossible if metaphorical content depended on illocutionary features”.

communicated by a face-value reading of their thought or belief (e.g. ‘the subject is the cream of my coffee’) in order to more easily bring out how the subject plays a role to the thinker that is somehow structured along the lines of cream in coffee. As in the interpersonal spoken case, thinking metaphorically in this way might draw the thinker’s attention to salient similarities between the role that the subject plays in the thinker’s life and the role that cream plays in a cup of coffee, or enable the thinker to exploit common associations with cream in coffee. Thought and belief, metaphorical or otherwise, is not constrained by conversational maxims – but when the thinker wishes to utter sentences to communicate the proposition expressed by her thought or beliefs, the utterances are so constrained. In the same way, metaphorical sentences in thought are not constrained by conversational maxims, and play the role of aiding the thinker in her practical and theoretical reasoning. But when the proposition expressed by the metaphorical sentence is uttered as part of a conversation, then the speaker gets the hearer to retrieve the speakers’ intended meaning by getting the hearer to engage in an inferential processes aimed at retrieving the non-face-value meaning⁹⁰.

Here, then, is how a Gricean metaphor fictionalist explains contextual variations in felicity of beliefs and thoughts which communicate propositions of the form “S knows that P”. In the low stakes case, Hannah comes to believe a true non-face-value proposition (e.g. that she has sufficient evidence in support of the proposition that the bank will be open tomorrow) from some other albeit false at face-value proposition (e.g. I know that the bank will be open tomorrow). Her belief is felicitous in the context because the proposition communicated by the non-face-value reading of her belief is true. In the high-stakes case, the proposition communicated by the non-face-value reading of her belief is false, and so the belief is not felicitous. The same principle applies to knowledge denials. In the low-stakes case, the proposition communicated by the non-face-value reading of “I don’t know that the bank will be open tomorrow” is infelicitous because the proposition communicated by a false non-face-value reading is false. Metaphor fictionalism, then, explains how knowledge talk can be used in thought to communicate its non-face-value content⁹¹.

Furthermore, the function of metaphor appears to cohere better with the function of knowledge talk than does the function of Gricean implicature in general. In typical cases of implicature, speakers choose to communicate their intended meaning via implicature because they wish to avoid being reported as having asserted the *implicated* content (Camp 2006). In the famous letter writing case, the referee chooses to communicate via implicature precisely because she does not wish to overtly commit herself to the propositional content ‘the student is a poor philosopher’. Given this, we might wonder why speakers would communicate content regarding epistemic positions via implicature. After all, knowledge attributions are not typically the types of speech acts that speakers wish to distance themselves from. That speakers do overtly and obviously

⁹⁰ Thank you to David Austin for helping me to understand how Gricean metaphor theorists treat metaphor in thought.

⁹¹ This is not to endorse a Gricean picture of metaphor as the correct account of metaphor: I leave the question of which picture of metaphor is correct open. The point is rather that if one is otherwise persuaded by a Gricean SPI picture, then further fleshing the view out so that the underlying pragmatic mechanism is that posited by Gricean metaphor confers advantages over the ‘bare’ Gricean SPI view.

commit to the communicated content is precisely what enables knowledge talk to serve its functions of recommending informants and ending inquiry.

Of course, this is not the only function of implicature and it is possible that knowledge talk is an example of implicature for purposes other than politeness and a desire to avoid overt commitment. What might this be? A possibility already discussed is §1 of this chapter that knowledge talk is communicated via implicature for emphasis. Another perhaps more plausible option is that using implicature increases verbal efficiency: one of the uses of implicature is to enable the communication of multiple contents through the utterance of a single content, thus avoiding the laborious task of asserting each content that they wish to communicate (see Levinson 2000: 28–31; Camp 2006: 3; Wilson and Sperber 2012: 62). For example, one might assert the single content “there is a petrol station around the corner” to implicate a multitude of contents, including “the petrol station is open”. However, this does not seem to apply to the case of knowledge talk. Following Meibauer (2009:378), Dinges (2016) distinguishes between additive implicature and substitutional implicature. Additive implicature conveys an implicated content in addition to the asserted content, as in the above example where implicated content ‘the petrol station is open’ is communicated in addition to the asserted content ‘there is a petrol station around the corner’. In contrast, substitutional implicature conveys the implicated content *instead of* the asserted content. For example, someone might sarcastically assert “John’s a fine friend” in order to implicate *instead* that John is a poor friend. Dinges argues, convincingly, that an SPI which posits that knowledge talk typically communicates via implicature must posit substitutional implicature. For speakers who utter sentences of the form “S knows that P” clearly do not intend to communicate *both* that S is infallible with regards to P, and additionally that S is in a good enough position with regards to P for S’s purposes. Therefore, it is implausible that on SPI knowledge talk serves the function of communicating multiple contents, of which the asserted content is one.

In contrast, metaphor fictionalism has a clear answer as to why the implicated metaphorical content would be communicated by uttering a sentence that is false on a face-value reading. The metaphor frame ‘S knows that P’ structures how we think about being in a good epistemic position with regards to P, making the high standards and the sense in which we can take a proposition as given salient. Metaphor fictionalism, even an implicature-based metaphor fictionalism, is therefore well placed to respond to the objection that implicature theories cannot explain how beliefs communicate non-face-value content.

6.4.2. *The awareness objection*

Kyriacou (2021:124) writes that critics of SPI often rely on a “single, often quickly aired” objection: the objection from semantic awareness. Roughly, the objection is that everyday speakers and hearers are not aware of using ‘know’ to communicate the content that SPI proposes. Given that the semantic intuitions of everyday speakers and hearers is generally a reliable indicator as to what speakers use their uttered words to mean, positing widespread semantic blindness counts against SPI. I will argue that although this objection does apply to ‘bare’ Gricean SPI, cashing out SPI in terms of metaphor fictionalism avoids the problem.

Following Kyriacou, we can distinguish between two variants of the semantic awareness objection. One strand objects to positing that speakers are unaware of the face-value semantic meaning of the term 'knows'. According to SPI, the face-value semantic meaning of the words 'knows' is that posited by the sceptic, something which everyday speakers and hearers appear unaware of. For example, Montminy (2009:652) writes that invariantism posits a "very puzzling type of error", namely, that "ordinary speakers lack access to the conventional meaning of knowledge sentences: although they are informed about all the relevant facts, they are unable to make a correct, literal use of such sentences". The second strand of the awareness objection is that in typical cases of implicature, speakers realise that there is a gap between the face-value meaning of the sentences that they use and the content that they intend to use those sentences to express, while speakers who engage in knowledge talk exhibit no such awareness. Dinges, for example, argues that SPI violates the principle that in cases of implicature "people are aware of the discrepancy between what they say and what they implicate" (2016:2588). This version of the objection is stronger in the sense that it does not hold that speakers who implicate are always able to spell out precisely what the face-value meaning of their sentences is: rather, it suffices that they are capable of detecting a gap between the face-value meaning of their utterance and the non-face-value meaning that they wish to express, even if they are unable to articulate precisely what the gap amounts to.

The awareness objection in both its iterations is reminiscent of the phenomenological objections to fictionalism discussed in Chapter 5, and I argued there that metaphor theory is well placed to answer these objections. In a nutshell, a lack of semantic awareness is to be expected on epistemic fictionalism, because it is typical to not notice the face-value semantic meanings of commonly used metaphors. Despite this, confrontational contexts do reveal what, following Yablo (2010), the fictionalists call 'felt gaps' between face-value and non-face-value content, and so metaphor theory also meets its version of the 'implicature awareness' objection. In what follows, I take a closer look at formulations of the awareness objections which target SPI, and argue that metaphor fictionalism has the resources to address them.

Let's start by considering the main piece of evidence that defenders of the awareness objection appeal to: intuitions regarding the felicity of responses in a series of dialogues. Specifically, defenders of the semantic awareness objection argue that in cases where a speaker communicates via implicature, hearers and speakers alike respond in ways that indicate that they do in fact recognise that they have used a form of implicature. In contrast, they argue that speakers who utter sentences of the form "S knows that P" do not exhibit such recognition.

Blome-Tillman takes exaggeration as a particular form of implicature, and asks us to compare the following two exchanges:

(1) **Exaggeration:**

A: I've been waiting forever!

B: That's absurd; no human being could ever wait that long.

A: Oh c'mon—I was just exaggerating!

(2) **Exaggeration K:**

H: I know that the bank will be open tomorrow.

S: That's absurd; no human being could ever know that.

H: Oh c'mon—I was just exaggerating⁹²

A's response to B seems entirely felicitous. Blome-Tillman argues that competent language users typically recognise that a face-value reading of A's first utterance would violate the maxim of Quality on account of it being transparently false, triggering the search for an implicated non-face-value reading. That A's first utterance is intended to implicate a non-face-value content is evidenced by the felicity of A's final utterance. In contrast, H's response does not strike us as felicitous. Blome-Tillman argues that if SPI were correct, H's final remark should strike hearers as being equally as felicitous as A's final remark. But unlike typical cases of implicature, speakers and hearers of sentences of the form "S knows that P" appear to be unaware of a face-value reading of the utterance, or indeed that there is a gap between what their utterance communicates on a face-value reading and what the speaker is using it to communicate.

Blome-Tillman argues that the same considerations rule out the metaphor version of SPI, as evidenced by the differences in felicity in the final line of the following two dialogues (2013:4301):

Metaphor:

A: She's made of stone, this girl.

B: That's absurd; human beings aren't made of stone.

A: Oh, c'mon—I was speaking metaphorically!

Metaphor K:

H: I know that the bank will be open tomorrow.

S: That's absurd; no human being could ever know that.

H: Oh, c'mon—I was speaking metaphorically!

I think that there are two related explanations for why H's response in 'Metaphor K' sounds infelicitous while A's response in 'Metaphor' sounds felicitous, and on neither is H's response infelicitous on account of being false. First, we should not expect speakers to identify what mechanism they are using to communicate their intended content. A's response is felicitous because B is obviously responding to the face-value reading of the metaphor, and in so doing, B's challenge makes it salient to hearers (A included) that A is speaking metaphorically. A expected B to respond to the non-face-value content that A wished to communicate to B, and

⁹² See also MacFarlane (2005:206).

prior to B's challenge had perhaps not even considered whether their utterance carries a face-value meaning. However, in making it obvious that they are responding to the face-value content, B makes the face-value reading salient to A, and forces A to either rephrase or recant – which here, A does by clarifying that they were speaking metaphorically. In contrast, S's challenge in 'Metaphor K' does not obviously – if at all – indicate that S is responding to any face-value reading of their original utterance. Without more context, it is unclear precisely what the nature of S's challenge is – they could just as plausibly be challenging the non-face-value content of H's original utterance. That is to say, S could be challenging whether given H's evidence (i.e. that H was at the bank two weeks ago on Saturday), any human being could be in a good enough position to act as an informant with regards to the proposition that the bank will be open tomorrow. Hence S's challenge does not make salient to H that H is speaking metaphorically, and further does not make it apparent to H that the cooperative conversational move is to rephrase or recant, for example by clarifying that H is speaking metaphorically.

A critic of metaphor fictionalism might respond that even if S's challenge were much more explicit – e.g. “that's absurd; no human being could ever know that the bank is open tomorrow, because no human being could rule out the possibility that they are a brain in a vat being deceived into believing that there is an external world containing banks to be open tomorrow” – H's reply of “Oh, c'mon—I was speaking metaphorically!” would still seem infelicitous. However, this is not a fatal objection to metaphor fictionalism. As discussed in chapters 3 and 4, many metaphorical sentences often are not registered as metaphorical by speakers or hearers, especially if the metaphors are conventional and familiar. Consider the following dialogue with a conventional metaphor:

Musical Notes

C: Her singing voice is amazing, she doesn't miss a note.

D: That's absurd; it isn't possible to hit or miss notes. She is vibrating her vocal cords at the correct frequency.

C: Oh, c'mon—I was speaking metaphorically!

D's response is highly uncooperative, but technically legitimate. Likewise C's response is correct, but might not sound entirely natural. Indeed, talk of hitting and missing notes might not strike us as metaphorical on first hearing; this is because this talk is standardly used to communicate its non-face-value meaning, such that it's perhaps hard to hear the face-value reading at all. On the picture defended in this thesis, whether a sentence is metaphorical is a function of whether it succeeds in communicating its implicated content in virtue of a structural mapping between the source domain and the target domain, irrespective of whether speakers are aware of such a mapping. Therefore, even when a fragment of discourse is metaphorical, the sentence “I was speaking metaphorically!” might not always sound natural.

This is not to say, however, that speakers are not on some level aware of a gap between the face-value readings of the sentences that they assert, and the non-face-value reading that they are using them to convey. Returning to the example where the brain in a vat challenge is made explicit, it

would perhaps be more natural for **H** to recant or rephrase, perhaps by protesting that she did not intend to rule out the **BIV** possibility. This suggests that **H** is aware of some gap between what her original utterance says if taken at face-value, and what she meant to communicate by uttering it, even if **H** is not able to identify precisely what linguistic mechanism is responsible for the gap. Metaphor fictionalism, then, avoids Blome-Tillman's version of the awareness objection to **SPI**⁹³.

This brings us on to the stronger formulation of the awareness objection. Targeting loose-talk versions of **SPI**, Dinges argues that in typical cases of implicature, speakers demonstrate awareness of a gap between what sentences assert and what speakers use them to communicate, while users of knowledge talk do not exhibit such awareness. As evidence, he contrasts the felicitousness of the final response in the following dialogue:

C'mon

A: It's three o'clock.

B: It's one minute past!

A: Oh, c'mon. You know what I meant.

Unlike Blome-Tillman's examples, the felicity of the 'Oh, c'mon. You know that I meant' response does not require that speakers be capable of articulating precisely what they meant, or via which mechanism. It only requires that speakers are aware that there is some difference between what they said and meant, and aware that their interlocutor has responded to the former

⁹³ I am not convinced that standard Gricean **SPI** falls foul of Blome-Tillman's objection, either; for Gricean implicature also often goes undetected by speakers in cooperative situations. For example, consider the following dialogue between the famous letter writer and a hypothetical admissions department:

Letter Writer

W: This student has excellent handwriting

A: We're not too concerned with handwriting quality, most assignments are typed nowadays anyway

W: Oh, c'mon—I was trying to implicate!

W's response also sounds unnatural. Perhaps this is accounted for by the inappropriateness of making an implicature explicit, but I think that speakers plausibly do not consciously partition their utterances according to whether they communicate content by implicature or by assertion. It seems that speakers using implicature are aware of what they are trying to communicate, and also aware of the desire to avoid committing too explicitly to what they are trying to communicate, rather than being explicitly aware of the fact that they are using implicature.

This is why confrontational contexts make salient the use of implicature – the admission's officer could have taken a confrontational approach and directly asked "What are you implying? You must know we don't care about that, are you trying to say that the student is unsuited to our programme?". This confrontational context makes salient the use of implicature, and forces the letter writer to either own or recant the implicated content of their utterance. Therefore, I agree with Dinges (2016) that whether speakers can identify the semantic content on a face-value reading of their utterance and further, whether they can identify the specific mechanism by which uttering a sentence succeeds in communicating their intended content is not an accurate test for ascertaining whether a given domain of discourse is typically used to implicate or assert.

as opposed to the latter. Indeed, in cases of implicature speakers can make sense of challenges that respond to the face-value content of the utterance, either due to genuine or deliberate misunderstanding, precisely because they are aware of the gap and so aware of the potential for misunderstanding. In the above dialogue, A is able to easily parse B's challenge as responding to the face-value content, which indicates that A is aware of a gap between the face-value content and the content that A wishes to communicate.

In contrast, Dinges argues that in the low-standards context, the final response in the following dialogue is infelicitous:

C'mon K

H: I know that the bank will be open.

S: I doubt that. The bank could have changed its hours.

H: Oh, c'mon. You know what I meant!

Dinges argues that sceptical pragmatic invariantism predicts that H's final utterance ought to be felicitous. If the final utterance were an implicature, then H should be aware of the gap between what she said and what she meant, and be able to easily parse that S is responding to what H said as opposed to what H meant. Yet, H's response is not felicitous. Therefore, SPI is to be rejected.

A natural response on behalf of the SPI proponent is that S *does* respond to what H meant as opposed to what H said. H's response, then, would be infelicitous because it would violate the maxim of Quality. Indeed, had S raised outlandish sceptical possibilities like the BIV possibility, the reply "Oh, c'mon. You know what I meant!" would appear entirely felicitous, suggesting that when there is a gap between speaker meaning and what is said regarding use of sentences of the form "S knows that P", agents are aware of it and respond appropriately. Perhaps, then, the mistake is to assume (as Dinges does) that the bank changing its hours is too unlikely an event to count as a serious response to what H intended to communicate with her first utterance in the low standards case. That is to say, the proponent of SPI can argue that this particular dialogue does not test for whether speakers can parse challenges aimed at what is said (as opposed to what is meant) because the challenge straightforwardly addresses what is meant.

In reply, Dinges argues that if the sceptical invariantist takes this route, she risks rendering SPI superfluous. For SPI appeals to implicature to explain the variability in intuitions regarding the assertability of sentences of the form "S knows that P" across low and high standard cases, such as the bank cases. According to SPI, what explains the variability is precisely that in the low-standards case, H does *not* intend to communicate that she can rule out the possibility that the bank has changed its hours. Therefore, the SPI proponent cannot easily claim that *in the low standards case* S is responding to what H meant, without undermining SPI's explanation for variability.

Further, Dinges argues that the proponent of SPI cannot claim that H's final response is acceptable, either, on pain of losing the sense in which higher standard contexts are supposed to

be enlightened contexts. As Hawthorne puts it, “as more and more possibilities of error become salient to us, we are reaching an ever more enlightened perspective” (2004:164–5). Cohen also holds that in a high standards context “one has a feeling of enlightenment regarding the correct application of the predicate in question” (2005:58). If this is right, then when S changes the context to one where more demanding epistemic standards are operative by making salient the possibility that the bank could have changed its hours, H ought to feel ‘enlightened’. But the response “Oh, c’mon. You know what I meant!” denies that H feels enlightened, and suggests that H had this possibility in mind all along.

I think that metaphor fictionalism versions of SPI can address both of these worries, and that whether H’s reply is felicitous depends on whether H takes S’s challenge to be a genuine attempt to respond to what H meant, or an uncooperative and deliberate misunderstanding. Take the case where S is being deliberately nitpicky. Perhaps it is part of the shared background assumptions between Hannah and Sara that the bank is extremely unlikely to change its hours, and so the challenge is responding to what is said as opposed to what is meant in a way that is obvious to both of them. In this case, “You know what I mean” is an appropriate response: Sara *did* know what Hannah meant, but chose to respond to what Hannah said, instead. Is this denying that making new possibilities salient induces feelings of enlightenment? No. But it does deny that Sara has created a more enlightened *context* by raising the possibility that the bank could have changed its hours. Rather, returning to Cohen’s phrasing, what Sara has done is enlighten hearers about the “correct application of the predicate in question”. For as metaphor theory predicts, when there is conflict about whether the face-value or non-face-value meaning should be taken up in a given conversational context, the face-value reading takes priority. In this example, Hannah may well feel enlightened as to which meaning wins out in this context, but need not feel that the context in general is a more enlightened one. This way of feeling enlightened passes Dinges’ awareness test: when challenged, Hannah responds in a way that demonstrates a gap between what she said and what she means. She is enlightened as to whether the face-value or non-face-value meaning is appropriate, when both are on the table. It is not a requirement of passing the test that Hannah were aware of the gap prior to making the claim and prior to being confronted with the challenge, back before the face-value reading was on the table. As with metaphor in general, speakers may not realise that there is a gap until the challenge is presented – hence the feeling of enlightenment.

Now, consider the case where Sara genuinely takes herself to be responding to what Hannah meant, as opposed to what she said. That is to say, Sara either believes that Hannah intended to communicate that her epistemic position is sufficient to rule out the possibility that the bank might change its hours, or Sara believes that their current context requires that it be so sufficient. The implicature-based metaphor fictionalism analysis of this version of ‘C’mon K’ can mirror that of the contextualist analysis given earlier in section 6.3.2.: Sara challenges the *aptness* of Hannah’s metaphorical sentence. Recall that whether “S know that P” is an apt metaphor for framing the epistemic position that S is in with regards to P is in part a function of how similar S’s epistemic position with regards to P is to infallibilism – and this similarity in turn is a function of the context. Here, Sarah pushes a dissimilarity between the source domain (infallible knowledge) and the target domain (Hannah’s epistemic position). Hannah can stick to her guns

by challenging whether the dissimilarity is sufficient to make the metaphor not apt (e.g. “Well, do you think that’s likely?”). But she cannot stick to her guns by saying “Oh, c’mon. You know what I meant” because Sarah has either actually succeeded in responding to what Hannah meant, and the reply is inappropriate; or Sarah does not know what Hannah meant, and the reply is infelicitous because clearly false. Crucially, this analysis does not detract from metaphor fictionalism’s analysis of the high-standards context. There, Sarah challenges what Hannah meant by challenging the aptness of Hannah’s metaphor; Hannah agrees that the metaphor is not apt; and Hannah retracts the implicated content⁹⁴.

This is not to say that non-metaphor based SPI cannot also meet these challenges, but that epistemic fictionalism based on an implicature metaphor theory has the resources to answer these objections should make it an attractive option to proponents of SPI. Those already convinced by SPI, then, might do well to adopt fictionalism as the way of making good on SPI.

6.5. Conclusion

In this chapter, I defended epistemic fictionalism by comparing it to rival views. I argued that epistemic fictionalism fares well against three antirealist views: hyperbole, loose talk and regulative ideals. In particular, understanding sentences of the form “S knows that P” as metaphorical better coheres with the function of knowledge talk, as well as linguistic data on how each type of talk behaves.

In discussing epistemic contextualism and sceptical pragmatic invariantism, I compared each view to a fictionalism ground in its ‘counterpart’ way of understanding metaphor. That is to say, I compared epistemic contextualism to epistemic fictionalism underpinned by a contextualist view of metaphor, and I compared sceptical pragmatic invariantism to epistemic fictionalism underpinned by a gricean implicature view of metaphor. In doing so, I hope to have shown two things: first, that epistemic fictionalism does not commit one to any particular view of metaphor, and leaves open the question of how sentences of the form “S knows that P” communicate the non-face-value content; second, that going fictionalist about knowledge talk brings something to the table beyond the ‘bare’ context-sensitivity or ‘bare’ implicature already posited by epistemic contextualism and sceptical pragmatic invariantism respectively. On the epistemic contextualist side, epistemic fictionalism claims that knowledge talk behaves like a metaphor because it *is* metaphorical, and thus is an improvement on contextualist views that posit that knowledge talk is merely analogous to a particular type of context-sensitive term. Furthermore, epistemic fictionalism has a principled explanation for why the sceptic succeeds in imposing her own reading of sentences of the form “S knows that P”, and matches the contextualist explanation of everyday context shifts. On the sceptical pragmatic invariantist side, epistemic fictionalism provides an explanation for why knowledge attributions in thought succeed in conveying a non-face-value content, and for responses under challenge.

⁹⁴ On the version of SPI that Dinges discusses and objects to, the face-value reading of ‘knows’ is the demanding but nevertheless attainable one operative in the high standards bank case. The metaphor fictionalist who adopts the same face-value meaning posited by Dinges has an even more straightforward explanation for the difference between the high standard and low standard cases, namely, that in the high-standard case Sarah responds both to what Hannah meant and to what she said.

I have limited the discussion to these two views because, like epistemic fictionalism, they posit that the sceptical infallibilist is right (at least in one sense) about the face-value reading of sentences of the form “S knows that P”, but also that such sentences are typically used to communicate other, non-face-value contents. The epistemic contextualist posits that both the sceptical infallibilist reading and the communicated content are legitimate face-value readings, while the sceptical pragmatic invariantist posits that the former is a face-value reading, and the latter a non-face-value reading. In showing how the epistemic fictionalist has the resources to address objections to these views while retaining their advantages, I hope to have gone some way towards demonstrating that epistemic fictionalism deserves further investigation.

Concluding Remarks

This thesis has outlined and defended a novel view in epistemology: epistemic fictionalism, according to which knowledge talk is metaphorical. Understanding knowledge talk as metaphorical is to understand the concept of knowledge as being underpinned by a metaphorical frame which maps elements of the source domain (knowledge on an infallibilist picture) onto a target domain (related to the functions of knowledge talk). According to my epistemic fictionalism, a face-value reading of sentences of the form “S knows that P” is that S is infallible with regards to P. However, such sentences are not typically used to express their face-value content, but to communicate their metaphorical content, roughly: that S is in a good enough position with regards to p to act as a reliable informant as to whether P, that S may terminate inquiry into whether P, and so on.

The claim that knowledge talk is metaphorical might initially strike readers as surprising and, perhaps, as *prima facie* implausible. My primary aim throughout this thesis has been to motivate epistemic fictionalism as a serious contender worthy of further development and investigation. I will now briefly summarise how each chapter contributes towards motivating epistemic fictionalism. To start, Chapter 1 canvassed some of the arguments in favour of infallibilism about knowledge, and diagnosed epistemologists’ resistance to infallibilism as being due infallibilism entailing high-standards scepticism. One of the key motivations for epistemic fictionalism is its ability to capture the sense in which sceptical infallibilism is correct, at least as it pertains to one reading of sentences of the form “S knows that P”. Chapter 2 addressed the New Apraxia objection to sceptical infallibilism, and argued that a concept or discourse domain may be functional in spite of typically being used in false sentences, thus motivating a broader class of anti-realist views to which fictionalism belongs.

Chapters 3 introduced fictionalism and laid out five criteria for a well-motivated fictionalist view, before applying these criteria to inform what epistemic fictionalism might look like, in particular, what content sentences of the form “S knows that P” might communicate on a face-value and non-face-value reading. Chapter 4 developed my positive proposal, a ‘function-first’ epistemic fictionalism underpinned by the mechanism of metaphor frames. It motivated epistemic fictionalism by spelling out the practical advantages of knowledge talk being framed in terms of infallibilism, and by highlighting how fictionalism retains the advantages of standard sceptical infallibilism while avoiding the damaging consequences of scepticism. On the infallibilist side, that the face-value reading of sentences of the form “S knows that P” is that posited by the infallibilist explains why infallibilism in general, and sceptical conclusions in particular, are compelling: they are, literally, true. Yet, epistemic fictionalism avoids the pitfalls commonly associated with sceptical infallibilism: sentences of the form “S knows that P” typically communicate a true content, and so everyday speakers are not committed to widespread error.

Chapter 5 moved on to defending epistemic fictionalism by addressing three pressing objections, including the phenomenological objection. Chapter 6 compared fictionalism to rival views, including other anti-realist views (hyperbole, loose talk and regulative ideals), epistemic contextualism and sceptical pragmatic invariantism. I have argued that epistemic fictionalism fares well against these views, both in terms of accounting for the function of knowledge talk and in explaining how knowledge talk behaves.

As noted, I hope to have shown that epistemic fictionalism is a promising view worthy of further development. The picture that I have presented here is amenable to being further specified and developed in different directions. The general claim that knowledge is metaphorical is compatible with a range of different theories on what the face-value and non-face-value readings of sentences of the form “S knows that P”. In this thesis, I argued that the face-value reading is that posited by the sceptical infallibilist, and the non-face-value reading is closely tied to the function of the concept of knowledge. But I have remained impartial on the question of what functions the concept of knowledge serves – any theory of the function of knowledge that is compatible with it fulfilling the functions that I have outlined may be plugged into my fictionalist theory. How this detail is filled in will directly bear on what content knowledge talk typically expresses. As such, epistemic fictionalism is compatible with a range of views on what knowledge talk typically communicates.

Epistemic fictionalism is also compatible with a wide range of metaphor theories. Whether one opts for a content or force, assertoric or non-assertoric, contextualist or implicature-based picture of metaphor will radically impact one’s theory of the mechanics which operate knowledge talk, and land one with different types of epistemic fictionalism. I think that this gets the order of investigation right: rather than picking the underlying account of her posited fictional discourse that best coheres with her target discourse domain, the job of the fictionalist is to establish whether a given domain of discourse does belong to a given type of fictional discourse, for example, metaphor. The details of the fictionalist view will then be beholden to how that entire domain of talk functions. My job in this thesis has been to make plausible the claim that knowledge talk is metaphorical talk – what exactly this means for the mechanics of knowledge talk is a further question to be answered by appeal to metaphor theory more generally.

Another aim of this thesis has been to showcase how the methods of function-first epistemology can fruitfully be put to work in developing an anti-realist theory. The focus on function has played multiple complementary roles, and entered the picture at different levels of analysis. In this sense, the scope of function-first epistemology has been pushed beyond investigating a concept’s intension, extension and value. To highlight this point, I will conclude by briefly pointing to how considerations of function have informed different aspects of the arguments in this thesis, starting with the micro level and working outwards.

First, in constructing the details of my epistemic fictionalist view, considerations of function directly inform the content that I take to be communicated by knowledge talk. On epistemic fictionalism, sentences of the form “S knows that P” communicate that S is a reliable informant with regards to P, that S may end inquiry into whether P, and so on. In this sense, the concept of knowledge is directly shaped by the function that knowledge talk exists to serve (Chapter 4).

The role of function goes beyond this, however. I have argued that keeping an eye trained to the function of domains of discourse as a whole illuminates what type of discourse they are. Here, I have argued from the claim that knowledge talk serves the function typical of metaphorical talk (specifically, of presenting information in cognitively accessible ways) to the conclusion that we have good reason to think that knowledge talk is in fact metaphorical. In particular, we have reason to think that knowledge talk is metaphorical, as opposed to being hyperbolic, being an instance of loose talk or reflective of a regulative ideals (Chapter 5). Considerations regarding the function of metaphorical talk in general not only make plausible the claim that knowledge talk is metaphorical, but also further vindicate hypotheses about the content that knowledge talk communicates. I argued from the claim that my particular construal of epistemic fictionalism (on

which sentences of the form “S knows that P” communicate function-related contents) enables knowledge talk to serve the function typical of metaphor talk, to the conclusion that we have good reason to suppose that these are indeed the contents communicated, as opposed to those proposed by Chung (2018) (Chapter 3). In this way, reflecting on the functions served by typical examples of a type of talk sheds light both on what type of talk the target discourse is, and on what it must be like (for example, what content it must express) to serve the functions typical of this type of talk.

On a wider scale, function plays a pivotal role in the construal of fictionalism itself. I argued that a well-motivated fictionalist theory needs some explanation of the function of uttering sentences that are typically infelicitous if read at face-value, and that the notion of fiction must be part of the explanation of how the discourse domain serves this function. On this picture, that a discourse domain be functional - and specifically, that it be functional *in virtue of* being fictional - is a requirement of fictionalism in general. This way of thinking about fictionalism bakes in the central function-first principle that a concept has the shape that enables it to fulfil its function, and expands the principle to say that a discourse domain comes equipped with the underlying mechanism that enables it to fulfil its function, whether this mechanism be a face-value or non-face-value mechanism.

Often arguments from function that yield an anti-realist theory such as fictionalism are classed as *debunking* arguments, showing that we have good reason to reject a domain of discourse⁹⁵. In contrast to both these debunking arguments and to the New Apraxia arguments (Chapter 2), I have argued that considerations of function can vindicate an anti-realist picture of knowledge. That is to say, I have argued that starting from considerations of function makes plausible the view that knowledge talk is underpinned by the mechanism of metaphor. This is no debunking project: that knowledge is functional, and functional *because* fictional, gives us every reason to keep on using our functional, fictional, infallibilist concept of knowledge.

⁹⁵ For example, function-led debunking arguments have been levelled at religion (Hobbes 1651, Hume 1757, Nietzsche 1887); gender (Haslanger 2012); morality (Mackie 1977, Joyce 2001), metaphysics (Ladyman and Ross, 2007); social order (Marx, 1867); to name just a few. Although cf: Queloz 2021.

Bibliography

- Adler, J. (2002). *Belief's own ethics*. Boston, MIT Press.
- Aikin, S. (2020). Does Metaphilosophically Pragmatist Anti-Skepticism Work? *Logos and Episteme*, 11(3), 391-398.
- Aldridge, M., & Luchjenbroers, J. (2007). Linguistic Manipulations in Legal Discourse: Framing questions and 'smuggling' information. *International Journal of Speech, Language and the Law*, 14(1), 85-107.
- Anderson, C. (2014). Fallibilism and the flexibility of epistemic modals. *Philosophical Studies*, 167(3), 597-606.
- Armour-Garb, B. & Woodbridge, J. (2015) *Pretense and Pathology: Philosophical Fictionalism and its Applications*. Cambridge: Cambridge University Press
- Armour Garb, B. & Kroon, F. (2020). Introduction. Armour-Garb, B. & Kroon, F. (Eds) *Fictionalism in Philosophy*, Oxford, Oxford University Press, 1-27
- Armstrong, D. (1986). The Nature of Possibility. *Canadian Journal of Philosophy*, 16(4), 575-594
- Aristotle. *Rhetoric*. In Roberts, W. R. (Trans).
- Arseneault, M. (2021). An implicature account of metaphorical perspective. *International Review of Pragmatics*, 13(1), 14-28.
- Austin, J. (1946). Other minds. *Aristotelian Society Supplementary Volume*, 20, 148- 87.
- Bach, K. (2005). The Emperor's New 'Knows' In G. Preyer and G. Peter (Eds). *Contextualism in Philosophy: Knowledge, Meaning and Truth*, 51-89.
- Balaguer, M. (1998). *Platonism and anti-Platonism in mathematics*. Oxford, Oxford University Press.
- Balaguer, M. (2009). Fictionalism, Theft, and the Story of Mathematics. *Philosophia Mathematica*, 17(2), 131-162.
- Barnden, J A. (2007). Metaphor, Semantic Preferences and Context-Sensitivity. In Ahmad, K., Brewster, C., Stevenson, M. (Eds). *Words and Intelligence II. Text, Speech and Language Technology*, vol 36, Springer, Dordrecht.
- Baumann, P. (2011). WAMs: Why worry? *Philosophical Papers*, 40(2), 155-177.
- Beebe, J. R. (2012). *Social Functions of Knowledge Attributions*. Oxford, Oxford University Press.
- Beebe, H. (2018). Philosophical Scepticism and the Aims of Philosophy. *Proceedings of the Aristotelian Society*, 118(1), 1-24.

- Bezuidenhout, A. (2001). Metaphor and what is said: A defense of a direct expression view of metaphor. *Midwest Studies in Philosophy*, 25(1), 156–186.
- Black, M. (1962). Metaphor. *Proceedings of the Aristotelian Society*, 55, 273-94
- Black, T. (2005). Classic Invariantism, Relevance, and Warranted Assertability Manœuvres. *Philosophical Quarterly*, 55, 328-336
- Blackburn, S. (2005). Quasi-Realism No Fictionalism. In Kalderon, M.E. (Ed.), *Fictionalism in Metaphysics*. Oxford, Oxford University press, 322–338
- Blome-Tillmann, M. (2008). The Indexicality of “Knowledge”. *Philosophical Studies*. 138(1), 29–53.
- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118(470), 241-294
- Blome-Tillmann, M. (2013). Knowledge and implicatures. *Synthese*, 190(18), 4293–4319.
- Bonjour, L. (2010). The Myth of Knowledge. *Philosophical Perspectives*, 24, 57–83.
- Booth, A. R. (2018). Advice for Infallibilists: DIVORCE and RETREAT!. *Synthese*. 195, 3773–3789
- Borgman, C.L., (1999). The user’s mental model of an information retrieval system: an experiment on a prototype. *International Journal of Human-Computer Studies*, 51, 435–452
- Bougher, L.D. (2012). The Case for Metaphor in Political Reasoning and Cognition. *Political Psychology*, 33, 145-163.
- Bourne, C. & Caddick Bourne, E. (2020). Folk Stories: What Has Fiction to Do with Mental Fictionalism?’ in Armour-Garb and Kroon (Eds.) *Fictionalism in Philosophy*, Oxford, Oxford University Press, 168-186
- Brandom, R. (2001). Reason, Expression, and the Philosophical Enterprise. In Ragland, C. P. & Heift, S. (Eds). *What is Philosophy?;* London: Yale University Press.
- Brennan, J. (2010). Scepticism about Philosophy. *Ratio*, 23(1), 1–16
- Brock, S. (2002). Fictionalism about Fictional Characters. *Noûs*, 36, 1–21.
- Brock, S. (2014). The Phenomenological Objection to Fictionalism. *Philosophy and Phenomenological Research*, 88, 574–592.
- Brown, J. (2006). Contextualism and Warranted Assertibility Manœuvres. *Philosophical Studies*, 130, 407-435
- Brown, J. (2018). *Fallibilism: Evidence and Knowledge*. Oxford: Oxford University Press
- Brueckner, A. (1994). The structure of the skeptical argument. *Philosophy and Phenomenological Research*, 54(4), 827-835
- Bueno, O. (2009). Mathematical Fictionalism. In O. Bueno and Ø. Linnebo (Eds.). *New Waves in Philosophy of Mathematics*, Hampshire, Palgrave Macmillan, 59–79.

- Burgess, J. P. (1983). Why I am not a nominalist. *Notre Dame Journal of Formal Logic*, 24(1), 93-105.
- Burgess, A. & Burgess, J. P. (2011). *Truth*. Princeton, Princeton University Press
- Butchvarov, P. (1970). *The concept of Knowledge*. Evanston, Northwestern University Press
- Caddick Bourne, E. (2018). Fictionalism in metaphysics. *Routledge Encyclopedia of Philosophy*, Taylor and Francis. Available at: <https://www.rep.routledge.com/articles/thematic/fictionalism-in-metaphysics/v-1/sections/1-fictionalisms-rationale-and-targets>
- Cameron L., Low G. & Maslen R. (2010). Finding systematicity in metaphor use. In Cameron L., Maslen R. (Eds). *Metaphor Analysis: Research Practice in Applied Linguistics, Social Sciences and the Humanities*. Equinox, pp. 116-46.
- Camp, E. (2006a). Contextualism, metaphor, and what is said. *Mind and Language*, 21(3), 280-309.
- Camp, E. (2006b). Metaphor and that certain 'je ne sais quoi'. *Philosophical Studies*, 129(1), 1 - 25.
- Camp, E. (2007). Showing, Telling and Seeing. Metaphor and “Poetic” Language. *Baltic International Yearbook of Cognition, Logic and Communication*, 3, 1-24
- Camp, E. (2009). Two Varieties of Literary Imagination: Metaphor, Fiction, and Thought Experiments. *Midwest Studies in Philosophy*, 33(1), 107-130.
- Camp, E. (2017). Why metaphors make good insults: perspectives, presupposition, and pragmatics. *Philosophical Studies*, 174, 47-64
- Camp, E. (2019). Perspectives and Frames in Pursuit of Ultimate Understanding. In Grimm, S.R. (ed.), *Varieties of Understanding: New Perspectives from Philosophy, Psychology, and Theology*, New York, Oxford Academic, pp. 17-46
- Cappelen, H. (2005) Pluralistic skepticism: Advertisement for speech act pluralism. *Philosophical Perspectives*, 19(1), 15-39.
- Carter, A. (1995). Knowledge and hyperbole. *Heythrop Journal*, 36(1), 46-64.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Carston, R., & Wearing, C. (2011). Metaphor, hyperbole and simile: A pragmatic approach. *Language and Cognition*, 3(2), 283-312
- Carston, R., & Wearing, C. (2015). Hyperbolic language and its relation to metaphor and irony. *Journal of Pragmatics*, 79, 79-92.
- Chalmers, David J. (2015). Why Isn't There More Progress in Philosophy?, *Philosophy* 90(1), 3-31
- Chisholm, R. M. (1966). *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall

- Chrisman, M. (2008). A dilemma for moral fictionalism. *Philosophical Books*, 49, 4-13.
- Chung, J. (2018). Could knowledge talk be largely non-literal? *Episteme*, 15(4), 383-411.
- Chung, J. (2020). Skeptical Arguments, Conceptual Metaphors and Cross-Cultural Challenges, in Mizumoto, M., Ganeri, J. & Goddard, C. (Eds) *Ethno-Epistemology: New Directions for Global Epistemology*. New York, Routledge, pp. 154-183
- Citron F.M. & Goldberg A.E. (2014). Metaphorical sentences are more emotionally engaging than their literal counterparts. *Journal of Cognitive Neuroscience*, 26(11), 2585-95.
- Claridge, C. (2011). *Hyperbole in English: a Corpus-Based Study of Exaggeration*, Cambridge, Cambridge University Press
- Clarke-Doane, J. (2012) Morality and Mathematics: The Evolutionary Challenge, *Ethics*, 122(2), 313-340
- Climenhaga, N. (2021). A Cumulative Case Argument for Infallibilism. In Christos Kyriacou & Kevin Wallbridge (Eds.), *Skeptical Invariantism Reconsidered*. Routledge.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*. 13, 57-89.
- Cohen, S. (2000a). Contextualism and Skepticism. *Philosophical Issues*, 10, 94-107.
- Cohen, S. (2000b). Replies. *Noûs*, 34, 132-9.
- Colston, H.L. (1997) Salting a wound or sugaring a pill: The pragmatic functions of ironic criticism. *Discourse Processes*, 23, 1, 25-45
- Colston, H. L., & Keller, S. B. (1998). You'll never believe this: Irony and hyperbole in expressing surprise. *Journal of Psycholinguistic Research*, 27(4), 499-513.
- Colston, H.L. & O'Brien, J. (2000) Contrast of Kind Versus Contrast of Magnitude: The Pragmatic Accomplishments of Irony and Hyperbole, *Discourse Processes*, 30:2, 179-199
- Colyvan, M. (2001). *The Indispensability of Mathematics*. Oxford, Oxford University Press.
- Conee, E. (2005). Contextualism contested. In M. Steup & E. Sosa (Eds.). *Contemporary Debates in Epistemology*. Malden, Blackwell Publishing. 47-56
- Copeland, B. J. (2002). The Genesis of Possible Worlds Semantics. *Journal of Philosophical Logic*, 31(2), 99-137.
- Cotterill, J. (1998). 'If it doesn't fit, you must acquit': metaphor and the O.J. Simpson criminal trial. *International Journal of Speech, Language and the Law*, 5(2), 141-158.
- Craig, E. (1987). XII*—The Practical Explication of Knowledge. *Proceedings of the Aristotelian Society*, 87(1), 211-226.
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford, Clarendon Press.

- Crimmins, M. (1998). Hesperus and Phosphorus: sense, pretense, and reference. *Philosophical Review*, 107, 1-47.
- Cuneo, T. (2006). Saying what we mean: an argument against expressivism. In Shafer-Landau, R. (Ed.), *Oxford Studies in Metaethics: Volume 1*. Oxford, Clarendon Press, pp. 35-71
- Cummings, E. E. (1932). Impressions IV. *Selected Poems 1923-1958*. Faber & Faber.
- Cunning, D. (2009). *Argument and Persuasion in Descartes' Meditations*. New York, Oxford University Press.
- Daly, C.J. (2006). Mathematical fictionalism - no comedy of errors, *Analysis*, 66(3), 208-216,
- Daly, C.J. (2008). Fictionalism and the attitudes. *Philosophical Studies*, 139, 423-440
- D'Amato, M. (2013). Buddhist Fictionalism. *Sophia*, 52 (3), 409-424.
- Davidson, D. (1978). What metaphors mean. *Critical Inquiry*, 5(1), 31-47.
- Davis, W. A. (2004). Are Knowledge Claims Indexical? *Erkenntnis*, 61, (2-3), 257-281.
- Davis, W. A. (2005). Contextualist theories of knowledge. *Acta Analytica*, 20(1), 29-42.
- Davis, W. A. (2007). Knowledge claims and context: loose use. *Philosophical Studies*, 132, 395-438
- DeRose, K. (1991). Epistemic Possibilities. *The Philosophical Review*, 100(4), 581-605
- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52, 913-929.
- DeRose, K. (1995). Solving the Skeptical Problem. *The Philosophical Review* 104, 1-52.
- DeRose, K. (2002). Assertion, Knowledge, and Context. *The Philosophical Review*, 111(2), 167-203.
- DeRose, K. (2004). Single Scoreboard Semantics. *Philosophical Studies*, 119, 1-21
- Descartes, R. (1996). *Meditations on First Philosophy: with selection from objections and Replies*. (Cottingham, J. Trans.). Cambridge, Cambridge University Press. (Original work published 1641)
- Dietrich, E. (2011). There Is No Progress in Philosophy. *Essays in Philosophy*, 12(2), 329-344.
- Dinges, A. (2016). Skeptical pragmatic invariantism: good, but not good enough. *Synthese*, 193(8), 2577-2593.
- Dinges, A. (2021). Knowledge and loose talk. In Kyriacou, C. and Wallbridge, K. (Eds) *Skeptical Invariantism Reconsidered*, London, Routledge, 272-297.
- Dodd, D. (2011). Against Fallibilism. *Australasian Journal of Philosophy*, 89 (4): 665-85.

- Dougherty, T. (2011). Fallibilism. In Duncan Pritchard & Sven Bernecker (Eds.), *The Routledge Companion to Epistemology*. Routledge.
- Dougherty, T. & Rysiew, P. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research* 78 (1):123-132.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67(24), 1007-1023.
- Dutant, J. (2007). The case for infallibilism. In C. Penco, M. Vignolo, V. Ottonelli, & C. Amoretti (Eds.), *Proceedings of the 4th latin meeting in analytic philosophy*, Genoa: University of Genoa. pp. 59-84
- Eklund, M. (2019). Fictionalism. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/fall2011/entries/fictionalism/>.
- Elmore, K. C., & Luna-Lucero, M. (2017). Light bulbs or seeds? How metaphors for ideas influence judgments about genius. *Social Psychological and Personality Science*, 8(2), 200-208.
- Emmet, D. (1994). *The Role of the Unrealisable: A Study in Regulative Ideals*, London: The MacMillan Press
- Engel, P. (2008). In What Sense Is Knowledge the Norm of Assertion?. *Grazer Philosophische Studien*, 77(1), 99-113
- Entman R. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4), 51-8
- Epley N, Waytz A, & Cacioppo JT. (2007). On seeing human: a three-factor theory of anthropomorphism. *Psychological Review*, 114(4), 864-866.
- Eshleman, A. S. (2005). Can an Atheist Believe in God? *Religious Studies*, 41(2), 183 - 199.
- Ervas, F. (2021). Metaphor, ignorance and the sentiment of (ir)rationality. *Synthese*, 198, 6789-6813
- Ervas, F., & Ledda, A. (2014). Metaphors in quaternio terminorum comprehension. *Isonomia*, 4, 179-202
- Everett, T. J. (2006). Antiskeptical Conditionals. *Philosophy and Phenomenological Research*, 73 (3), 505-536
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111(1), 67-94.
- Fantl, J. & McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford, Oxford University Press.
- Feldman, R. (1997). Review Essay: Knowledge and the Stated of Nature: An Essay in Conceptual Synthesis. *Philosophy and Phenomenological Research*, 57(1), 205-221.
- Feldman, R. (1999). Contextualism and Skepticism. *Philosophical Perspectives*, 13, 91-114.
- Field, H. (1989). *Realism, Mathematics and Modality*. Oxford, Blackwell

- Fillmore C. J. (1975). An alternative to checklist theories of meaning. *Proceedings of the First Annual Meeting of the Berkeley Linguistics Society*, Berkeley, 123–31.
- Fine, R. (1993). Fictionalism. *Midwest Studies In Philosophy*, 18(1), 1–18.
- Fratantonio, G. (2021). No Infelicity for the Infallibilist. *Theoria*, 87, 1307-1321
- Fricker, M. (2008). Scepticism and the genealogy of knowledge: situating epistemology in time. *Philosophical Papers*, 37(1), 27-50.
- Fricker, M. (2016). What's the point of blame? A paradigm-based explanation. *Noûs*, 50(1), 165-183.
- Friend, S. (2008). Hermeneutic moral fictionalism as an anti-realist strategy. *Philosophical Books*, 49(1), 14-22
- Friend, S. (2012). Fiction as a Genre. *Proceedings of the Aristotelian Society*, 112, 179–209
- Frigg, R. (2010). Models and Fiction. *Synthese*, 172(2), 251-68.
- Foucault, M. (1971). *Nietzsche, genealogy, history*. In *Language, counter-memory, practice selected essays and interviews by Michel Foucault* (Bouchard, D. F., Ed.) Ithaca, NY: Cornell University Press, 139–64.
- Fumerton, R. (1990). Metaepistemology and Skepticism. In: Roth, M.D. & Ross, G. (Eds) *Doubting. Philosophical Studies Series*, 48, Springer, Dordrecht.
- Fumerton, R. (2006). *Epistemology*, Oxford, Blackwell
- Gardiner, G. (2015). Teleologies and the methodology of epistemology. In J. Greco & D. Henderson (Eds.), *Epistemic Evaluation: Purposeful Epistemology*. Oxford: Oxford University Press.
- Gatia, D. E. (2010). Colour Fictionalism. *Rivista di Estetica*, 43, 109-123
- Gawthorne, D. (2013). Fictionalising Jurisprudence: An Introduction to Strong Legal Fictionalism. *Australian Journal of Legal Philosophy*, 38, 52–73.
- Geary, J. (2011). *I is an other*. New York: HarperCollins
- Gelfert, A. (2011). Steps to an Ecology of Knowledge: Continuity and Change in the Genealogy of Knowledge, *Episteme*, 8(1), 67-82.
- Gettier, E. L. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121-123.
- Gibbs, R.W. Jr., Buchalter, D.L, Moise, J.F. & Farrar, W.T. (1993). Literal meaning and figurative language. *Discourse Processes*. 16(4), 387-403
- Gibbs, R.W. Jr., & Colston, H. L. (2012). *Interpreting figurative meaning*. Cambridge, Cambridge University Press.
- Gildea, P. & Glucksberg, S. (1983). On understanding metaphor: the role of context. *Journal of Verbal Learning and Verbal Behavior*, 22(5), 577-590.

- Glucksberg, S., Gildea, P. & Bookin, H A. (1982). On understanding nonliteral speech: Can people ignore metaphors?. *Journal of Verbal Learning and Verbal Behavior* 21(1), 85-98
- Goddard, C. (2010). Universals and variation in the lexicon of mental state concepts. In B. Malt & P. Wolff (Eds.), *Words and the mind*. Oxford, Oxford University Press.
- Godfrey-smith, P. (2009). Models and fictions in science. *Philosophical Studies*, 143(1), 101-116.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73(2), 771-791.
- Graham, P. (2015). Epistemic Normativity and Social Norms. In David Henderson & John Greco (Eds.), *Epistemic Evaluation: Point and Purpose in Epistemology*. Oxford, Oxford University Press. pp. 224-244.
- Greco, J. (2008). What's Wrong with Contextualism? *The Philosophical Quarterly*, 58(232), 416-436.
- Grice, H.P. (1975). Logic and Conversation. In P. Cole, & J. L. Morgan. (Eds.), *Syntax and Semantics*, 3, New York: Academic Press. pp. 41-58.
- Grimm, S. (2015). Knowledge, Practical Interests, and Rising Tides. In J. Greco & D. Henderson (Eds.), *Epistemic Evaluation*, Oxford, Oxford University Press.
- Halliday, D. (2007). Contextualism, comparatives and gradability. *Philosophical Studies* 132, 381-393
- Hannon, M. (2013). The Practical Origins of Epistemic Contextualism. *Erkenntnis*, 78(4), 899-919.
- Hannon, M. (2014) Fallibilism and the Value of Knowledge. *Synthese*, 191, 1119-1146.
- Hannon, M. (2015). The universal core of knowledge. *Synthese*, 192, 769-786.
- Hannon, M. (2019a). *What's the Point of Knowledge?: A Function-first Epistemology*. New York: Oxford Scholarship Online
- Hannon, M. (2019b). Skepticism: Impractical, Therefore Implausible. *Philosophical Issues*, 29, 143-158
- Hannon, M. (2022). Book Review: The Practical Origins of Ideas by Matthieu Queloz. *Mind*. Available at: <https://academic-oup-com.ezproxy.is.ed.ac.uk/mind/advance-article/doi/10.1093/mind/fzac004/6537097>
- Hannon, & Nguyen, J. (forthcoming). Understanding philosophy. *Inquiry*. 1-37.
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford, Oxford University Press.
- Haspelmath, M., & Tadmor, U. (2009). World loanword database. Max Planck Digital Library. <http://wold.livingsources.org/>.

- Hauser, D. & Schwarz, N. (2019). The War on Prevention II: Battle Metaphors Undermine Cancer Treatment and Prevention and Do Not Increase Vigilance. *Health Communication*, 35(13), 1698-1704.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford, England: Oxford University Press.
- Hawthorne, J. & Stanley, J. (2008) Knowledge and action. *Journal of Philosophy*, 105, 571–590.
- Henderson, D. (2009). Motivated Contextualism. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 142(1), 119–131.
- Henderson, D. (2011). Gate-Keeping Contextualism. *Episteme*, 8(1), 83-98.
- Henderson, D. & Greco, J. (2015). Introduction: the Point and Purpose of Epistemic Evaluation. In Henderson, D. & Greco, J. (Eds.). *Epistemic Evaluation: Purposeful Epistemology*. Oxford, Oxford University Press. pp. 1-28
- Hills, D. (1997). Aptness and truth in verbal metaphor. *Philosophical Topics*, 25, 117–153
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press.
- Hobbes, T. (2017). *Leviathan*. (Brooke, C. Ed.). Oxford, Clarendon Press (Original work published 1651)
- Hookway, C. (2002). Emotions and epistemic evaluations. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The Cognitive Basis of Science* Cambridge: Cambridge University Press. pp. 251-262.
- Horwich, P. (2012). *Wittgenstein's Metaphilosophy*. Oxford, UK: Oxford University Press
- Huemer, M. (2001). *Skepticism and the Veil of Perception*, Oxford, Rowman and Littlefield Publishers
- Huemer, M. (2007). Moore's paradox and the norm of belief. In S. Nuccetelli & G. Seay (Eds.), *Themes from G.E. Moore* (pp. 142–57). Oxford, Oxford University Press.
- Hume, D. (2014). Of the origin of justice and property. *A Treatise of Human Nature*. (Selby-Bigge, L.A. & Nidditch, P.H. Eds). Second edition. Oxford, Oxford University Press (Original work published 1739)
- Hume, D. (2014). *The Natural History of Religion*. (Colver, A. W. & Price, J. V., Eds.). Oxford, Oxford University Press (Original work published 1757)
- Jackson, A. (2012). Two ways to put knowledge first. *Australasian Journal of Philosophy*, 90(2), 353–369.
- Jaster, R. (2013). Contextualism and Gradability. *GAP.8 Proceedings*. 318-323
- Jenkins, C. S. (2006). Lewis and Blackburn on quasi-realism and fictionalism. *Analysis*, 66(4), 315-319

- Johnson, M. (2010). Metaphor and Cognition. In Schmicking, D., Gallagher, S. (Eds) *Handbook of Phenomenology and Cognitive Science*. Springer, Dordrecht.
- Joyce, R. (2001) *The Myth of Morality*. Cambridge University Press.
- Joyce, R. (2005). Moral Fictionalism. In Mark Eli Kalderon (Ed.), *Fictionalism in Metaphysics*. Oxford, Oxford University Press. pp. 287-313.
- Joyce, R. (2013). Psychological Fictionalism, and the Threat of Fictionalist Suicide. *The Monist*, 96(4), 517–538.
- Joyce, R. (2020). Fictionalism: morality and metaphor. In Armour-Garb, B. & Kroon, F. (Eds) *Fictionalism in Philosophy*, Oxford, Oxford University Press, 103-121
- Kalderon, M.E. (2005). *Moral Fictionalism*, Oxford, Oxford University Press
- Kaplan, M. (2008). Austin's Way with Skepticism. In Kaplan, M. and Greco, J. (Eds) *The Oxford Handbook of Skepticism*. Oxford, Oxford University Press. pp. 348-371
- Kappel, K. (2010). On Saying That Someone Knows: Themes from Craig. In A. Haddock, A. Millar, A., & D. Pritchard (Eds.), *Social Epistemology*. Oxford, Oxford University Press
- Kelp, C. (2011). What's the Point of "Knowledge" Anyway?. *Episteme*. Edinburgh, Edinburgh University Press, 8(1), 53–66
- Kornblith, H. (2000). The Contextualist Evasion of Epistemology. *Philosophical Issues*, 10, 24–32.
- Kornblith, H. (2002). *Knowledge and Its Place in Nature*. Oxford, Oxford University Press.
- Kratzer, A. (1977). What *Must* and *Can* Must and Can Mean. *Linguistics and Philosophy*, 1(3), 337–55
- Kratzer, A. (1981). The Notional Category of Modality. In Eikmeyer, H.J. & Rieser, H. (Eds.), *Words, Worlds, and Contexts: New Approaches in Word Semantics*, Berlin: de Gruyter. pp. 38–74.
- Kratzer, A. (1991). Modality. In von Stechow, A. & Wunderlich, D. (Eds.), *Semantics: An International Handbook of Contemporary Research*, Berlin: de Gruyter, pp. 639–50.
- Kroon, F. (2001). Fictionalism and the Informativeness of Identity. *Philosophical Studies*, 106, 197–225.
- Kroon, F. (2011), Fictionalism in Metaphysics. *Philosophy Compass*, 6, 786-803.
- Kvanig, J.L. (2003) *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Kusch, M. (2009). Testimony and the Value of Knowledge. In Pritchard, D. Haddock, A. Millar, A. (Eds). *Epistemic Value*, Oxford, Oxford University Press, 60-94
- Kusch, M. & McKenna, R. (2020). The Genealogical Method in Epistemology'. *Synthese*, 197, 1057-1076

- Kyriacou, C. (2020). Assertion and Practical Reasoning, Fallibilism and Pragmatic Skepticism. *Acta Analytica*, 35(4), 543-561.
- Kyriacou, C. (2021). Semantic Awareness for Skeptical Pragmatic Invariantism. *Episteme*, 18(2), 123-141.
- Ladyman, J., & Ross, D. (2007). *Every Thing Must Go : Metaphysics Naturalized*. Oxford, Oxford University Press.
- Lakoff, G. & Johnson, M. (1980). *Metaphors We Live By*. Chicago, Chicago University Press
- Landau, M.J., Keefer, L, & Rothschild, Z.. (2014). Epistemic Motives Moderate the Effect of Metaphoric Framing on Attitudes. *Journal of experimental social psychology*, 53, 125–138.
- Landau M.J., Meier B.P. & Keefer L.A. (2010). A metaphor-enriched social cognition. *Psychological Bulletin*, 136(6), 1045-67.
- Lasersohn, P. (1999). Pragmatic Halos. *Language*, 75(3), 522–551.
- Lawlor, K. (2013). *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford, Oxford University Press.
- Le Poidevin, R. (2020). Fiction and the Agnostic. *European Journal for Philosophy of Religion*, 12(3), 163-181
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. Boston, The MIT Press.
- Levy, A. (2011). Information in Biology: A Fictionalist Account. *Noûs*, 45(4), 640-657
- Lewis, D. (1978). Truth in Fiction. *American Philosophical Quarterly*, 15, 37–46.
- Lewis, D. (1991). *Parts of Classes*. Oxford, Basil Blackwell
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Liggins D. (2010). The Autism Objection to Pretence Theories. *The Philosophical Quarterly*, 60, 764-782.
- Liggins, D. (2012). Fictionalism. *Oxford Bibliographies in Philosophy*. Available at: <https://www.oxfordbibliographies.com/display/document/obo-9780195396577/obo-9780195396577-0034.xml>
- Liggins, D. (2020). Against hermeneutic fictionalism. In Armour-Garb, B. & Kroon, F. (Eds.), *Fictionalism in Philosophy*. Oxford, Oxford University Press. 81-102
- Littlejohn, C. (2011). Concessive Knowledge Attributions and Fallibilism. *Philosophy and Phenomenological Research*, 83(3), 603-619.
- Littlejohn, C. (2013). The Russellian retreat. *Proceedings of the Aristotelian Society*, 113, 293–320.
- Logue, H. (2022). Gender Fictionalism, *Ergo*, 8, 125–162.

- Ludlow, P. (2005). Contextualism and the new linguistic turn in epistemology. In Gerhard Preyer & Georg Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford, Oxford University Press. pp. 11–51.
- MacFarlane, J. (2005). The Assessment Sensitivity of Knowledge Attributions. In Gendler, T. & Hawthorne, J. (Eds.) *Oxford Studies in Epistemology*, Oxford, OUP: 197-233.
- Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*, Harmondsworth, New York: Penguin.
- Marcus, R. (2015). *Autonomy Platonism and the Indispensability Argument*, Lanham, Rowman and Littlefield.
- Martinich, A. (1984). A theory for metaphor. *Journal of Literary Semantics*, 13, 35 – 56
- Marx, K. (1976). *Capital: A Critique of Political Economy. Volume One*. (Fowkes, B. Trans.). London, Penguin Books
- Mayer, R.E. (1976). Some conditions of meaningful learning for computer programming: advanced organizers and subject control of frame order. *Journal of Educational Psychology*, 68, 143–150
- McCarthy, & Carter, R. (2004). “There’s millions of them”: hyperbole in everyday conversation. *Journal of Pragmatics*, 36(2), 149–184.
- McElree B, & Nordlie J. (1999). Literal and figurative interpretations are computed in equal time. *Psychonomic Bulletin and Review*, 6(3), 486-94.
- McGlynn, A. (2019). Review of Jessica Brown, Fallibilism: Evidence and Knowledge. *Notre Dame Philosophical Reviews*. Available at: <https://ndpr.nd.edu/news/fallibilism-evidence-and-knowledge>
- McGrath, M. (2015). Two Purposes of Knowledge-Attribution and the Contextualism Debate. In David Henderson & John Greco (Eds.), *Epistemic Evaluation: Point and Purpose in Epistemology*. Oxford University Press, pp. 138-157
- McKenna, R. (2015). Epistemic contextualism defended. *Synthese*, 192, 363–383
- McKenna, R.(2023). *Non-Ideal Epistemology*, Oxford, Oxford University Press
- Meibauer, J. (2009). Implicature. In Mey, J. (Ed.), *Concise Encyclopedia of Pragmatics*, Amsterdam, Elsevier, 365–378
- Mehta, N. (2016). Knowledge and other norms for assertion, action, and belief: A teleological account. *Philosophy and Phenomenological Research*, 93(3), 681–705.
- Menzies, P. & Pettit, P. (1994). In Defence of Fictionalism about Possible Worlds. *Analysis*, 54(1), 27–36
- Miller, K. (2021). Pretence fictionalism about the non-present. *Inquiry*. 1-35
- Montminy, M. (2009). Contextualism, Invariantism and Semantic Blindness. *Australasian Journal of Philosophy*, 87(4), 639-657.

- Moore, G.E. (1959). *Philosophical papers*. New York: Macmillan
- Nagel, J. (2010). Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives*, 24 (1), 407-435.
- Nagel, J. (2011). The Psychological Basis of the Harman-Vogel Paradox. *Philosophers' Imprint*, 11, 1-28.
- Neta, R. (2011). A Refutation of Cartesian Fallibilism. *Noûs*, 45, 658-695.
- Newman, S. (1999). Uses of metaphor in Legal Argument. *New York Law Journal* (col. 3)
- Newton, L. & O'Sullivan, A. (2023). A Function-First Approach to Risk. *Unpublished Manuscript*
- Nietzsche, F. (1998). *On the Genealogy of Morality* (Clark, M. & Swensen, A.J., Trans.). Indianapolis, Hackett. Abbrev. (Original work published in 1887).
- Nogales, P. (2012). Metaphorical content as what is said. *Journal of Pragmatics*, 44(8), 997-1008
- Nolan, D. (2005). Fictionalist Attitudes about Fictional Matters. In Kalderon, M. E. (Ed.), *Fictionalism in Metaphysics*. Oxford, Clarendon Press. pp. 204-233.
- Nolan, D. (2016). Modal Fictionalism. *Stanford Encyclopedia of Philosophy*. Online, available at: <https://plato.stanford.edu/ENTRIES/fictionalism-modal/>
- Nolan, D., Restall, G. & West, C. (2005) Moral fictionalism versus the rest. *Australasian Journal of Philosophy*, 83(3), 307-330
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Pasnau, R. (2017). *After Certainty: A History of Our Epistemic Ideals and Illusions*. Oxford: Oxford University Press.
- Pritchard, D. (2001). The Opacity of Knowledge. *Essays in Philosophy*, 2 (1), 1-17.
- Pritchard, D. (2002) Recent Work on Radical Skepticism. *American Philosophical Quarterly*, 39, 215-257.
- Pritchard, D. (2005a). *Epistemic luck*. Oxford University Press UK.
- Pritchard, D. (2005b). The structure of sceptical arguments. *Philosophical Quarterly*, 55(218), 37-52.
- Pritchard, D. (2008). Radical Scepticism, Epistemic Luck, and Epistemic Value. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 82, 19-41.
- Pritchard, D. (2011). Epistemological disjunctivism and the basis problem. *Philosophical Issues*, 21(1), 434-455.
- Pritchard, D. (2012). Anti-Luck Virtue Epistemology. *The Journal of Philosophy*, 109(3), 247-279.

- Pritchard, D. (2015). Risk. *Metaphilosophy*, 46, 436-461.
- Pritchard, D. (2018). Contextualism and radical scepticism. *Synthese*, 195, 4733-4750
- Popa-Wyatt, M. (2020). 'Mind the Gap: Expressing affect with hyperbole and hyperbolic compounds.' In Barnden, J. & Gargett, A. (Eds). *Producing Figurative Expression: Theoretical, experimental and practical perspectives*. Amsterdam: John Benjamins Publishing Company. 449-468.
- Queloz, M. (2021). *The Practical Origin of Ideas*. Oxford: Oxford University Press.
- Rapp, A., Leube, D., Erb, M., Grodd, W. & Kircher, T. (2007). Laterality in metaphor processing: Lack of evidence from functional magnetic resonance imaging for the right hemisphere theory. *Brain and language*, 100, 142-9.
- Recanati, F. (2004). *Literal Meaning*. Cambridge University Press, Cambridge
- Reed, B. (2002). How to Think about Fallibilism. *Philosophical Studies*, 107, 143-157
- Reed, B. (2010). A Defense of Stable Invariantism. *Noûs*, 44 (2), 224-244
- Reed, B. (2012). Fallibilism. *Philosophy Compass*, 7(9), 585-596
- Reed, B. (2013). Fallibilism, epistemic possibility and epistemic agency. *Philosophical Issues*, 23(1), 40-69.
- Reimer, M. (1996). The Problem of Dead Metaphors. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 82(1), 13-25.
- Reimer, M. (2001). Davidson on metaphor. *Midwest Studies in Philosophy*, 25(1), 142-155.
- Reisfield G. & Wilson G., (2004). Use of metaphor in the discourse on cancer. *Journal of Clinical Oncology*, 22(19) 4024-7
- Rescher, N. (1987) *Ethical Idealism: An Inquiry into the Nature and Function of Ideals*, Berkley and Los Angeles: University of California Press
- Reuchamps, M., Dodeigne, J., & Perrez, J. (2018). Changing your political mind: The impact of a metaphor on citizens. *Regional & Federal Studies*, 28(2), 151-175
- Reynolds, S. L. (2002). Testimony, Knowledge, and Epistemic Goals. *Philosophical Studies* 110, 139-61.
- Reynolds, S. L. (2009). Making up the truth. *Pacific Philosophical Quarterly*, 90(3), 315-335
- Rinard, S. (2021) 'Pragmatic Skepticism', *Philosophy and Phenomenological Research*, 00:1-20.
- Ritchie L. D. (2013). *Metaphor*. Cambridge, Cambridge University Press
- Robins, S. & Mayer, R. (2000). The Metaphor Framing Effect: Metaphorical Reasoning About Text-Based Dilemmas. *Discourse Processes*, 30, 57 - 86.

- Rosen, G. (1990). Modal Fictionalism. *Mind*, 99, 327–354.
- Rosen, G. & Dorr, C. (2002). Composition as a fiction. In Gale, R. (Ed.), *The Blackwell Companion to Metaphysics*, Oxford, Blackwell. pp. 151–174.
- Rouse, J. (2008). Laboratory Fictions. In Fiction in Science. In Suárez, M. (Ed) *Philosophical Essays on Modeling and Idealization*. Routledge. pp. 37–55
- Rubio-Fernández, P. Wearing, C. & Carston, R. (2015). Metaphor and Hyperbole: Testing the Continuity Hypothesis, *Metaphor and Symbol*, 30(1), 24-40,
- Russell, G.K. (2021). Fancy loose talk about knowledge. *Inquiry*, 1-32.
- Rysiew, P. (2001). The Context-Sensitivity of Knowledge Attributions. *Noûs*, 35(4), 477–514.
- Rysiew, P. (2007). Speaking of Knowing. *Noûs*, 41, 627-662.
- Rysiew, P. (2012). Epistemic scorekeeping. In Brown, J. & Gerken, M. (Eds.), *Knowledge Ascriptions*. Oxford, Oxford University Press.
- Ruiz de Mendoza, F.J. & Galera, A. (2014). *Cognitive modeling. A Linguistic Perspective*, Amsterdam, John Benjamins.
- Sainsbury, M. (2010), *Fiction and Fictionalism*, London: Routledge.
- Sanford, D. (1991). Proper Knowledge. In McLaughlin, B.P. (Ed.), *Dretske and his Critics*, Oxford: Oxford University Press: 38-5 1
- Schaffer, J. (2004). Skepticism, Contextualism, and Discrimination. *Philosophy and Phenomenological Research*, 69, 138-155.
- Schaffer, J. (2008). Knowledge in the Image of Assertion. *Philosophical Issues*, 18, 1-19.
- Schiffer, S. (1996). Contextualist Solutions to Scepticism. *Proceedings of the Aristotelian Society*, 96, 317-333.
- Searle, J. (1969). *Speech Acts*, Cambridge, Cambridge University Press
- Semino E. (2008). *Metaphor in Discourse*. Cambridge, Cambridge University Press
- Semino, E. Demjén, Z. & Demmen, J. (2018). An Integrated Approach to Metaphor and Framing in Cognition, Discourse, and Practice, with an Application to Metaphors for Cancer, *Applied Linguistics*, 39(5), 625–645
- Siegel, H. (1997). *Rationality Redeemed?* Routledge.
- Shapiro, S. (1997). *Philosophy of Mathematics: Structure and Ontology*. Oxford, Oxford University Press.
- Smith, M. R. (2007). Levels of Metaphor in Persuasive Legal Writing. *Mercer Law Review*, 58(3)

- Smith, M. (2022). Underdetermination and closure: Thoughts on two sceptical arguments. In Jope, M. & Pritchard, D. (Eds.), *New Perspectives on Epistemic Closure*, Routledge.
- Smithies, D. (2015). Why Justification Matters. In Henderson, D. & Greco, J. (Eds.), *Epistemic Evaluation: Point and Purpose in Epistemology*. Oxford University Press. 224-244.
- Sperber, D. & Wilson, D. (1986). *Relevance: Communication and cognition*. (2nd ed.), Oxford: Blackwell.
- Sperber, D. & Wilson, D. (1995). Postface. In *Relevance: Communication and cognition*. (2nd ed.), Oxford: Blackwell, 255-279.
- Sperber, D. & Wilson, D. (2008). A deflationary account of metaphors. In Gibbs, R. W. (Eds.), *The Cambridge handbook of metaphor and thought*, Cambridge: Cambridge University Press, 84-105.
- Sprevak, M. (2013). Fictionalism about Neural Representations, *The Monist*, 96(4), 539-560
- Stanley, J. (2001). Hermeneutic Fictionalism. *Midwest Studies in Philosophy*, 25, 36-71
- Stanley, J. (2005) *Knowledge and Practical Interests*, Oxford, OUP.
- Stanley, J. (2008). Knowledge and Certainty. *Philosophical Issues*, 18, 35-57.
- Steen, G. J., Dorst, A. G., Herrmann, J. B., Kaal, A., Krennmayr, T., & Pasma, T. (2010). *A method for linguistic metaphor identification: From MIP to MIPVU*. Amsterdam, John Benjamins
- Stern, J. (1985). Metaphor As Demonstrative. *The Journal of Philosophy*, 82(12), 677-710
- Stern, J. (2000). *Metaphor in Context*. Cambridge, MA : MIT Press
- Stone, J. (2000). Skepticism as a Theory of Knowledge. *Philosophy and Phenomenological Research*, 60(3), 527-545.
- Stoutenburg, G. (2021). Skeptical Invariantism, Considered. In Kyriacou, C. & Wallbridge, K. (Eds.), *Skeptical Invariantism Reconsidered*. 80-101. Routledge.
- Stroud, B. (1984). *The Significance of Philosophical Scepticism*. Oxford, Oxford University Press.
- Thibodeau P.H. & Boroditsy, L. (2011). Metaphors we think with: The role of metaphor in reasoning. *PLoS ONE*, 6(2), 1-11
- Thibodeau, P.H. & Boroditsky, L. (2013). Natural Language Metaphors Covertly Influence Reasoning', *PLoS ONE*, 8(1), e52961
- Thibodeau, P.H. & Boroditsky, L. (2015). Measuring Effects of Metaphor in a Dynamic Opinion Landscape'. *PLoS ONE*, 10(7), 1-22
- Thibodeau, P. H., Matlock, T. & Flusberg, J. S. (2019). The role of metaphor in communication and thought. *Language and Linguistics Compass*, 13, 1-18

- Thompson, N. (2022). Setting the story straight: fictionalism about grounding. *Philosophical Studies*, 179, 343–361
- Thornburg, E. (2013). Metaphors Matter: How Images of Battle, Sports and Sex Shape the Adversary System. *Wisconsin Women's Law Journal*, 10, 225-281
- Toon, A. (2016). Fictionalism and the Folk. *The Monist*, 99, 280-2956
- Turri, J. (2016). Knowledge, certainty, and assertion. *Philosophical Psychology*, 29(2), 293-299
- Unger, P. (1971). A Defense of Skepticism. *The Philosophical Review*, 80(2), 198–219.
- Unger, P. (1975) *Ignorance : A Case for Scepticism*. Oxford: Clarendon Press.
- Vaihinger , H. (1924) *The Philosophy of 'As If' : A system of the theoretical, practical and religious fictions of mankind* (Ogden, C. K., Trans). London, Trubner & Co. (Original work published in 1911).
- Van Inwagen, P. (1977). Creatures of Fiction, *American Philosophical Quarterly*, 14(4).
- Van Frassen, B. (1980). *The Scientific Image*, Oxford University Press.
- Vazard, J. (2019). (Un)reasonable doubt as affective experience: obsessive–compulsive disorder, epistemic anxiety and the feeling of uncertainty. *Synthese*, 198(7), 6917-6934.
- Venturinha, N. (2019). Non-sceptical Infallibilism, *Analysis*, 80, 186-195.
- Walton, K. (1993). Metaphor and Prop-Oriented Make-Believe. *European Journal of Philosophy*, 1, 39–56.
- Walton, K. (2013). Metaphor, Fictionalism, Make-Believe: Response to Elisabeth Camp. Online: <https://philpapers.org/archive/WALPOM-3.pdf>
- Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114, 864 – 886.
- Wearing, C. (2006). Metaphor and what is said. *Mind and Language*, 21(3), 310–332
- Wearing, C. (2012). Metaphor, Idiom, and Pretense. *Noûs*, 46(3), 499-524.
- Wielenberg, E. J. (2010). On the Evolutionary Debunking of Morality. *Ethics*, 120(3), 441–464.
- Williams, B. 1978. *Descartes: The Project of Pure Inquiry*. Routledge.
- Williams, B. (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton, Princeton University Press.
- Williams, M. (2004). Scepticism and the Context of Philosophy. *Philosophical Issues*, 14, 456–475.
- Williamson, T. (2000) *Knowledge and its limits*. Oxford, Oxford University Press.

- Wright, C. (2005). Contextualism and Scepticism: Even-Handedness, Factivity and Surreptitiously Raising Standards. *The Philosophical Quarterly*, 55(219), 236–262
- Yablo, S. (1998). Does Ontology Rest on a Mistake? *Aristotelian Society Supplement*, 72: 229–6.
- Yablo, S. (2000). A paradox of existence. In T. Hofweber & A. Everett (eds.), *Empty Names, Fiction, and the Puzzles of Non-Existence*. Cambridge, Cambridge University Press. pp. 275–312.
- Yablo. (2010). *Things papers on objects, events, and properties*. Oxford, Oxford University Press.
- Yourgrau, P. (1983). Knowledge and Relevant Alternatives. *Synthese*, 55(2), 175-190