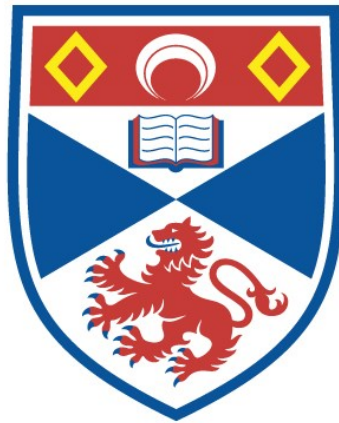


WHAT DO I KNOW?
SCEPTICISM: REASONING AND KNOWLEDGE

Laurence Carrick

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews



2018

Full metadata for this item is available in
St Andrews Research Repository
at:
<http://research-repository.st-andrews.ac.uk/>

Please use this identifier to cite or link to this item:
<http://hdl.handle.net/10023/13770>

This item is protected by original copyright

What do I know?

Scepticism: Reasoning and Knowledge.

By

Laurence Carrick

Department of Philosophy
University of St Andrews

Supervisors:

Dr. Philip Ebert
Dr. Patrick Greenough



University of
St Andrews

This thesis is submitted in partial fulfilment for the
degree of

DOCTOR OF PHILOSOPHY
at the University of St Andrews.

DATE OF SUBMISSION 25th January 2018

Candidate's declaration

I, Laurence Gregory Carrick, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 67,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree.

I was admitted as a research student at the University of St Andrews in January 2012.

I confirm that no funding was received for this work.

Date

Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date

02/05/2018

Signature of supervisor:

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Laurence Gregory Carrick, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date

Signature of candidate

Date

02/05/2018

Signature of supervisor:

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Laurence Gregory Carrick, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date

Signature of candidate

ACKNOWLEDGEMENTS

Firstly, I would like to thank my supervisors. I owe much to Philip Ebert for his tireless encouragement, enthusiasm, and—ultimately—patience. Without Philip’s willingness to entertain my inchoate ideas and unorthodox approaches, it is unlikely I would have ever realised my thesis. Likewise, I would like to thank Patrick Greenough for his immense support in writing this thesis. I would also like to thank Derek Ball and Justin Snedegar, respectively, for agreeing to be my examine me, and for giving me with invaluable feedback in my annual reviews.

Secondly, I thank my friends Joe Slater, Fenner Tanswell, Mark Bowker, Ste Broadrick, Elspeth Gillespie, and Carley Hollis. We all came to St Andrews in 2006, and have shared so many wonderful experiences together. I cannot imagine a better group of friends. Long may that continue. I extend special thanks to my friends Ravi Thakral and Paul Black for their generosity and love. I cannot express enough thanks for Morag and Ian Hamilton for everything, and above all else, the opportunity to escape from the delusions and illusions of philosophy. I thank Celia Hindley, Tigger and Tabitha for reminding me why I do it—and being a reliable source of cat related content.

Finally, My greatest thanks go to my mother. Without her this thesis would not have been possible. She taught me how to write to express my thoughts, but perhaps more importantly, she taught me (knowingly or not) the value of learning, and how to do it. This thesis stands as a testament to what I have learnt: when one is done asking questions, and challenging the orthodox answers, one is done with life.

ABSTRACT

This thesis concerns approaches to solving the problem of paradoxical sceptical arguments from ignorance within contemporary epistemology. In chapter 1, I critically discuss three frameworks for approaching the sceptical problem, and argue that theoretical responses are unsatisfactory. In chapter 2, I critically examine recent accounts of sceptical hypotheses, and argue against them on the grounds of generality, and in favour of my own account. In chapter 3, I critically examine recent accounts of the epistemic principles underwriting sceptical arguments from ignorance, and argue against them on the grounds of generality, and in favour of my own account. In chapter 4, I critically evaluate the adequacy of resolutions to sceptical paradoxes suggested by three prominent versions of epistemological contextualism. In chapter 5, I examine a central objection to the error theories implied by contextualist resolutions of sceptical paradoxes, which focuses on the notion of semantic blindness. Two assessments of the objection are set out, and contextualist responses to each. I argued that considerations of semantic blindness count against contextualist resolutions of sceptical paradoxes in favour of invariantists. In chapter 6, I assess the potential for an invariantist to provide an adequate error-theory concerning, and resolving, sceptical paradoxes. I critically assess approaches based on aspects of the heuristics and biases paradigm, and of dual-process theories of mindreading. I propose, instead, a novel anti-sceptical error-theory in terms of the default-interventionist model of dual-process theory of judgement and reasoning, together with my conclusions from chapters 2 and 3.

TABLE OF CONTENTS

Introduction.	vii
1 The Problem of Scepticism.	1
1.1 Scepticism.	1
1.2 Characterisations of Responses to Scepticism.	6
1.3 Theoretical and Therapeutic responses: Happy-face Solutions?	20
1.4 Concluding remarks.	30
2 Sceptical Hypotheses.	31
2.1 Introduction.	31
2.2 Sceptical Hypotheses and Sensitivity.	32
2.3 Experiential constraint.	35
2.4 The case against the experiential, and sensitivity constraints.	37
2.5 Indication constraint.	43
2.6 Bad possibilities.	46
2.7 Proposed condition on sceptical hypotheses.	51
2.8 Concluding remarks.	55
3 Epistemic Principles.	59
3.1 Introduction	59
3.2 Preliminaries.	60
3.3 Candidate epistemic principles.	61
3.4 Principle of Ignorance Possibilities.	66
3.5 Principle of Epistemic Possibility.	71
3.6 Concluding remarks.	78
4 Contextualism and Scepticism.	81
4.1 Introduction.	81
4.2 Epistemic Contextualism and Scepticism.	82

4.3	General Framework of a Contextualist Resolution.	84
4.4	Versions of Epistemic Contextualism and their Resolutions.	88
4.5	Concluding remarks.	107
5	Contextualist Error-Theories	109
5.1	Introduction.	109
5.2	Contextualism and Semantic Blindness.	110
5.3	Semantic Blindness: knockdown objection?	113
5.4	Semantic blindness: scorecard approach.	122
5.5	Concluding remarks.	127
6	Invariantist Error-Theories.	129
6.1	Introduction.	129
6.2	Semantic versus Psychological Error Theories.	130
6.3	Availability and Error.	131
6.4	‘Mindreading’ and Dual Process Theories.	135
6.5	Invariantism, Error and ‘Sceptical pressure’	137
6.6	Dual Processes and Default-Interventionism.	148
6.7	Default-interventionism, Sceptical Illusions, and Invariantism.	155
6.8	Assessing my proposal.	161
6.9	Concluding remarks.	164
	Bibliography	165

Introduction.

Each concept means what it singly means, and nothing else; and if the conceiver does not know whether he means this or means that, it shows that his concept is imperfectly formed.

— William James, (James, 1911, 48-9)

This thesis is concerned with the problem of paradoxical sceptical arguments from ignorance, and, subsequently, how it may be solved. I begin, in chapter 1, by setting the stage for the rest of the thesis. I start with some important preliminaries, in which I attempt to clarify precisely as possible the focus of the thesis: the problem of scepticism.

After setting out the problem in terms of sceptical arguments and their paradoxical nature, I move on to critically discuss and compare three different frameworks for thinking about how to solve the problem of sceptical paradoxes. In this discussion, I identify the two cruxes of the subsequent thesis: what constraints a fully satisfactory resolution of sceptical paradoxes must meet, and the diagnostic work necessary to begin assessing whether they are met. I take on the latter in the diagnostic work comprising chapters 2 and 3. In chapters 4 and 5 I take on the former in assessing a prominent attempt to resolve sceptical paradoxes. Finally, in chapter 6, I draw the chapters together to suggest a way to fully satisfactory resolution of sceptical paradoxes.

In majority of chapter 1, I focus on critically examining various characterisations of approaches to the providing a solution the problem of sceptical paradoxes. I examine three different ways of characterising solutions to sceptical problems: Pryor's (2000) distinction between *ambitious* and *modest* anti-sceptical projects, Schiffer's (2004) distinction between *happy*-, and *unhappy-face* resolutions to sceptical paradoxes, Williams' (1991) distinction between theoretical and *therapeutic* diagnoses of sceptical arguments. I argue, ultimately, that Pryor's distinction is not useful in seeking a solution to sceptical problems, and that the distinctions drawn by Schiffer and Williams together provide a more useful dimensions along which to find a solution to the sceptical problem. I develop an argument that theoretical diagnosis of sceptical arguments face second-order problems to providing happy-face resolutions of sceptical paradoxes. I apply the argument to reject Pryor's anti-sceptical project, and use this to motivate the subsequent search for a happy-face resolution that is compatible with what I term a confusion-therapeutic diagnosis of sceptical arguments.

In chapter 2, I focus on asking whether we really understand the role played by sceptical hypotheses in sceptical arguments. I critically assess and argue against various

prominent accounts of sceptical hypotheses. The discussion is guided by the methodological assumption that a certain generality is required from an adequate account of sceptical hypotheses. I introduce a novel example and draw on it to argue against various prominent accounts of sceptical hypotheses. Instead, I propose my own account of sceptical hypotheses, and argue that it is adequately general.

In chapter 3, I focus on asking whether we really understand what epistemic principle actually underwrite sceptical arguments. I critically assess and argue against various prominent proposals for this principle. Here, again, the discussion is guided by a methodological assumption that an account of the epistemic principle underwriting a sceptical argument which generalises is preferable to one that does not. I argue, on this basis, that that various accounts should be rejected. Instead, I consider a recently proposed account of the epistemic principle underwriting sceptical arguments that does generalise. Nonetheless, I raise doubts about the explanatory adequacy of this proposal. Finally, I make my own proposal for the principle underwriting sceptical arguments in terms of a novel analysis of epistemic possibility, and defend this account as both the sufficiently general, and adequately explanatory.

In chapter 4, I give a detailed critical discussion of the putative resolutions of sceptical paradoxes proffered by versions of epistemological contextualism in terms of the constraints on a fully satisfactory resolutions identified in chapter 1. Next, in chapter 5 I examine a central objection to the error theories of contextualist resolutions that focuses on the notion of semantic blindness. I discuss two lines of along which the objection can be interpreted, and the contextualist response to them. Here, I argue that considerations of semantic blindness give reason to think that contextualism is at a significant disadvantage to its competitor, invariantism, in terms of providing a fully satisfactory resolution of sceptical paradoxes.

In chapter 6, I assess the potential for the epistemological view of invariantism to provide an adequate error-theory concerning sceptical paradoxes. I critically assess approaches to explaining our erroneous judgments in terms of recent developments in cognitive psychology, including the heuristics and biases paradigm, and dual-process theories of mindreading. I argue that some putative candidates for an anti-sceptical invariantist error-theory given in terms of cognitive terms are inadequate. Instead, I put forward a novel anti-sceptical error-theory given in terms of the default-interventionist model of dual-process theory of judgement. I conclude that this provides a *prima facie* fully satisfactory solution to the sceptical problem.

THE PROBLEM OF SCEPTICISM.

1.1 Scepticism.

This thesis concerns itself with scepticism and the problem which it poses within epistemology. When writing about scepticism, Craig (1990, 104) suggests that it is prudent to begin by saying what exactly one is talking about. The task of describing what scepticism is, and the problem it presents, within epistemology, is not a straightforward one. This is because scepticism, as Stroud observes, “has been different things at different times in the history of philosophy, and has been put to different uses” (Sosa and Stroud, 1994, 87). For this reason, it is widely accepted that the term ‘scepticism’ does not denote a homogeneous set of issues. According to Williamson, for example, “[s]cepticism is a disease individuated by its symptoms (such as immodest protestations of ignorance); we should therefore not assume that it can be caused in only one way” (Williamson, 2000, 165). And, similarly, Wright advises caution in this respect:

Much literature on the topic of external world scepticism proceeds as if there were a single general form of problem that has to be confronted. In fact, however, arguments for epistemological scepticism come in a variety of significantly different forms and, while generality is of course a merit in a response, there is no reason to expect that they should succumb to a uniform treatment. (Wright, 1991, 87)

Wright, here, illuminates two issues when considering how to go about discussing the

problem of scepticism. The first is that ‘scepticism’ is best understood as being a thesis of some description about knowledge or justification. The second is that a sceptical thesis presents an epistemological problem in virtue of being argued for. In this chapter, I will aim to unpack that laconic statement of the topic of this thesis.

I begin, in sections 1.1.1 and 1.1.2, by setting out the sceptical thesis, the type of argument for it with which I am interested, and the problem posed by the paradoxical nature of these arguments. In section 1.2, I outline and discuss three distinct approaches to distinguishing ways on solving the problem of sceptical paradoxes. Here I examine and critically reject Pryor’s distinction between ambitious and modest responses to scepticism, and consider in detail both Schiffer’s distinction between happy-face and unhappy-face resolutions of sceptical paradoxes, as well as Williams’ distinction between theoretical and therapeutic diagnoses of sceptical arguments. In section 1.3, I draw together considerations of Schiffer and Williams’ distinctions to argue against the potential for theoretical solutions to provide the happy-face resolutions of sceptical paradoxes. I conclude that the most promising route to a happy-face resolution lies in the direction of a diagnosis of sceptical arguments that explicates our error regarding sceptical paradoxes in terms that are not purely epistemological.

1.1.1 Sceptical arguments and hypotheses.

In its most general sense, scepticism can be understood as the epistemological thesis “that nobody knows anything, or that nobody has good reason to believe anything” at all (Sosa and Stroud, 1994, 291). However, for my purposes here, I use the term scepticism in a more restricted sense to refer to the idea that *nobody knows, or has good reason to believe any* empirical proposition. So why is scepticism supposed to be interesting? The above quote from Wright highlights an important point that, insofar as scepticism presents a problem, it does so because there may be arguments put forward in support of it.¹ This thesis only concerns itself with sceptical arguments that share the following single common ‘argument from ignorance’ form; where *S* refers to some epistemic subject, *p* refers to some empirical proposition, and *h* refers to some ‘hypothesis’:

ARGUMENT FROM IGNORANCE.

- (i) *S* doesn’t know that $\neg h$.
- (ii) If *S* knows that *p*, then *S* knows that $\neg h$.

¹ Cf. Williams (2013).

So,

(iii) S doesn't know that p .²

Hereafter 'sceptical argument' will refer only to those which have the argument from ignorance, or AI, form. But not every instance of AI is sceptical. In order for an instance of AI to be a *sceptical* argument in the sense meant here, the argument needs to apply to a broad range of empirical propositions p . Whether it does this or not, will depend on the specific hypothesis h involved. The hypothesis that *my bicycle was just stolen*, for example, seems apt to be used in an AI argument to the effect that I do not know that my bike is where I left it. Consider the following argument:

BICYCLE.

(1) I do not know that my bicycle was not just stolen.

(2) If I know that my bicycle is where I left it, then I know it was not just stolen.

So,

(3) I do not know that my bicycle is where I left it.

But whilst the hypothesis that *my bicycle was just stolen* seems apt to raise a challenge to my claimed knowledge of my bicycle's whereabouts, it is not apt to challenge an obvious range of empirical propositions I might claim to know. However, this is precisely what the hypothesis must do in a sceptical argument, and only those that do so will be referred to as sceptical hypotheses. This can be illustrated by the following two examples of sceptical hypotheses that have featured prominently in the literature on scepticism: the *dreaming* hypothesis, and the *brain in a vat* hypothesis.

In its most basic form, the dreaming hypothesis is that one is currently dreaming. Its origins lie in the work of Rene Descartes's *Meditations*, where he writes:

How many times have I dreamt at night that I was in this place, dressed, by the fire, although I was quite naked in my bed? it certainly seems to me at the moment that I am not looking at this paper with my eyes closed; that this head that I shake is not asleep; that I hold out this hand intentionally and deliberately, and that I am aware of it. What happens in sleep does not seem

² The argument form is well known, and often presented in superficially different ways. I follow Brueckner (2011) in my presentation.

as clear and distinct as all this. But in thinking about it carefully, I recall having often been deceived in sleep by similar illusions, and, reflecting on this circumstance more closely, I see so clearly that there are no conclusive signs by which one can distinguish clearly between being awake and being asleep, that I am quite astonished by it; and my astonishment is such that it is almost capable of persuading me that I am asleep now. (Descartes, 1982, 96-7)

Here we have our first example of a sceptical hypothesis with Descartes' introduction of the possibility that he is dreaming. That is, a hypothesis which is apt to be used in a *sceptical* argument. Where dr is the hypothesis that S is currently dreaming, it seems that the following line of reasoning can be put forward, for every epistemic subject S and *any* empirical proposition p :

DREAMING.

(4) S doesn't know $\neg dr$.

(5) If S knows that p , then S knows that $\neg dr$.

So,

(6) S doesn't know that p .

The second sceptical hypothesis to consider is the *brain in a vat* hypothesis, which is, perhaps, the most commonly discussed example of a sceptical hypothesis.³ To get a sense of what this involves, it is instructive to consider the following quote from Hilary Putnam:

[I]magine that a human being (you can imagine this to be yourself) has been subjected to an operation by an evil scientist. The person's brain (your brain) has been removed from the body and placed in a vat of nutrients which keeps the brain alive. The nerve endings have been connected to a super-scientific computer which causes the person whose brain it is to have the illusion that everything is perfectly normal. There seem to be people, objects, the sky, etc; but really all the person (you) is experiencing is the result of electronic impulses traveling from the computer to the nerve endings. The

³ Its origins can be traced back to (Unger, 1975).

computer is so clever that if the person tries to raise his hand, the feedback from the computer will cause him to ‘see’ and ‘feel’ the hand being raised. Moreover, by varying the program, the evil scientist can cause the victim to ‘experience’ (or hallucinate) any situation or environment the evil scientist wishes. (Putnam, 1981, 5-6).

What sceptical import do these considerations have? To see this, first let *S* stand for some human being. Next let ‘*biv*’ stand for the proposition that *S* is in exactly the situation that Putnam asks you to imagine in the above passage—they are a *brain in a vat* (or, a BIV for short). Finally, let ‘*hands*’ stand for the proposition that *s has hands*. This gives us the second sceptical argument:

ENVATTED BRAIN.

(7) *S* does not know that $\neg biv$.

(8) If *S* knows that *hands*, then *S* knows that $\neg biv$.

So,

(9) *S* does not know that *hands*.

That concludes the set up with respect to the sorts of sceptical arguments with which I will be interested in discussing. In the next section, I will set out the problematic nature of these arguments.

1.1.2 Sceptical Paradoxes.

One way to characterise the problem of sceptical arguments is to argue that “sceptical arguments present us with classical sceptical paradoxes, and the problem of scepticism is the problem of solving those paradoxes” (Schiffer, 1996, 317). On taking this approach it is first pointed out that a classical paradox is “set of mutually inconsistent propositions each of which enjoys some plausibility when considered apart from the others” (Schiffer, 1996, 328). Next, it is argued that “[t]he premises of [a sceptical argument]... and the negation of its conclusion comprise such a set” (Schiffer, 2004, 178). This is because, it is argued, “each premise seems to some degree creditable when considered on its own, but the two together entail a conclusion we are apt to feel has got to be false” (Schiffer, 2004, 165). This characterisation of the sceptical problem maintains that the premises of arguments such as DREAMING and ENVATTED BRAIN strike us as *prima facie* plausible.

So let us examine the case for saying that the first premises of these sceptical arguments are plausible.

Many epistemologists do think that it is eminently plausible to deny both that we know *we are not currently dreaming*, or that *we are not brains in vats*. We often find descriptions of our epistemic position with respect to sceptical hypotheses in which these are characterised in the following way:

the very sort of propositions which one seems unable, in principle, to know are the denials of radical skeptical hypotheses, since these hypotheses concern scenarios which are phenomenologically indistinguishable from everyday life. (Pritchard, 2002, 217)

The idea here that sceptical arguments are paradoxical is in part due to the fact that sceptical hypotheses are in principle unknowable by anyone. They are designed to be such, so that anyone who might wonder whether or not they knew it to be false, would have no way to tell, since there would be nothing in the course of their actual experience that would be distinguishable from the hypothesis obtaining. Compare the situation with the non-sceptical BICYCLE argument and its first premise. In contrast to the DREAMING and ENVATTED BRAIN arguments, it is not *in principle* plausible to deny that I know that the *stolen* hypothesis is false. Rather, the question of whether I know that my bike has not been just stolen seems to be a contingent on the what I have done. It might seem implausible, for example, if I *have just this second* seen it where I left it. It might also be implausible to deny I know it has just been stolen if I have been able to see the only possible access to it for the entire time since I left it, and *not* seen anyone go by. In either case, the hypothesis that it has just been stolen would be distinguishable from my actual course of experience; I would have just seen it, or I would have (presumably) seen someone come by.

My interest, here, lies in those sceptical arguments from ignorance that are paradoxical in the sense I have described here. I now move on, to focus on considerations of how the problem of these paradoxes could be solved.

1.2 Characterisations of Responses to Scepticism.

In this section, I discuss how to solve the problem of scepticism by discussing various characterisations of approaches to the providing a solution. I examine three different ways of characterising solutions to sceptical problems. First, I discuss Pryor's distinction

between *ambitious* and *modest* anti-sceptical projects. I argue that this distinction between ambitious and modest responses to scepticism is not useful in seeking a solution to sceptical problems. Next, I outline then Schiffer's distinction between *happy*-, and *unhappy-face* solutions to sceptical problems. Lastly, I discuss Williams' distinction between theoretical and *therapeutic* diagnoses of sceptical problems. I suggest that the distinctions drawn by Schiffer and Williams together provide a more useful dimensions of assessment of putative solutions to the sceptical problem. In the final section, I discuss some ways in which Schiffer and William's distinctions relate to each other, apply these considerations to Pryor's work, and draw conclusions about how to best approach solving the sceptical problem.

1.2.1 Pryor's Anti-sceptical Projects.

There is a distinction concerning approaches to solving the problem posed by sceptical arguments that has been drawn by Pryor (2000). This is the distinction between what he terms 'ambitious' and 'modest' anti-sceptical projects. Firstly, I focus on setting out what the two anti-sceptical projects are supposed to be, and the attitudes Pryor takes towards their prospects for success. I offer a tentative critical assessment of the distinction before moving on, in the subsequent sections, to discuss relevant distinctions described by Schiffer (1996) and Williams (1991).

Pryor considers the possibility of responding to sceptical arguments, and, subsequently, solving the paradoxes which they present by way of what he calls an *ambitious anti-sceptical project*. He explains what this involves in the following way:

The ambitious anti-skeptical project is to refute the skeptic on his own terms, that is, to establish that we can justifiably believe and know such things as that there is a hand, using only premises that the skeptic allows us to use. (Pryor, 2000, 517)

The ambitious anti-sceptical project can be compared to what has been called the 'convince-the-sceptic' game. Our objective in this game is "to try to convince the sceptic that we have knowledge of the external world, without 'begging the question' against her. Specifically, we cannot appeal to any premise that is known on the basis of perceptual experience, because the legitimacy of such premises is precisely what the sceptic doubts" (Byrne, 2004, 301). In contrast to the ambitious anti-sceptical project of *refuting* sceptical arguments, Pryor considers an alternative possibility of responding to sceptical problems

by means of what he calls a *modest anti-sceptical project*. He explains what this involves in the following way:

The modest anti-skeptical project is to *establish to our satisfaction* that we *can* justifiably believe and know such things as that *there is a hand*, without contradicting obvious facts about perception. (Pryor, 2000, 517)

The modest and the ambitious anti-sceptical projects both succeed, only if they establish that the conclusion of sceptical arguments are false. This is because the conclusions of sceptical arguments is that we are unable to know, or have good reason to believe any empirical proposition—such as that *I have hands*, for example—, and *both* anti-sceptical projects aim to *establish* that we are able to do so. The contrast, then, between the ambitious and the modest anti-sceptical projects lies in that the ambitious project involves demonstrating that the sceptic *can't* have the premises they need for their argument (since it has been refuted). The modest anti-skeptical project, on the other hand, merely

attempts to diagnose and defuse those skeptical arguments; to show how to retain as many of our pretheoretical beliefs about perception as possible, *without accepting the premises the skeptic needs for his argument*. (Pryor, 2000, 517, my emphasis)

Pryor is not alone in distinguishing what he calls *ambitious* from *modest* 'anti-sceptical projects', and it is well appreciated that

many of the most influential responses to scepticism deny that an adequate response should answer sceptical doubts or use only principles which the sceptic would accept... [Such] accounts do not attempt to resolve doubt about whether one is a BIV or convince someone who doubts this that she has knowledge. Rather, they attempt to use part of our conception of the world to show how, despite the sceptical argument, we can have knowledge. (Brown, 2005, 7)

An approach to sceptical problems, for example, akin to what Pryor calls a modest anti-sceptical project, can be found in the work of Robert Nozick (1981). He explains that, regarding sceptical problems:

Our task here is to explain how knowledge is possible... In doing this, we need not convince the sceptic, and we may introduce explanatory hypotheses that he would reject. What is important for our task of explanation and understanding is that we find those hypotheses acceptable or plausible[.] (Nozick, 1981, 197-8)

Pryor is pessimistic when he states that the “prospects for this ambitious anti-skeptical project seem somewhat dim” (Pryor, 2000, 517). In fact, Pryor maintains that the prospects for ambitious anti-sceptical projects are even worse than this; he says “the ambitious anti-skeptical project is hopeless: we can’t demonstrate to the skeptic, using only premises he’ll accept, that we have any perceptual knowledge [or justification]” (Pryor, 2000, 520). In contrast, of course, Pryor is optimistic about the prospects of success for a modest anti-sceptical project, since it only “aims to set our own minds at ease, it’s not a condition for succeeding at it that we restrict ourselves to only making assumptions that the skeptic would accept” (Pryor, 2000, 518). This leads me to raise some general concerns about Pryor’s pattern of attitudes towards modest and ambitious anti-sceptical projects. In I have in mind, specifically, his optimism concerning the prospects of success for modest projects and his pessimism about those for ambitious anti-sceptical projects. And I will set out this concern next.

1.2.2 Doubts about Pryor’s Projects.

My concern with respect to Pryor’s distinction between ambitious and modest responses to sceptical arguments, can be articulated in what I take to be the following an uncontroversial point. If we are to be satisfied that we have established something, then we want to have *actually* established it. When aiming to establish something, we do not *just* want to do so using premises that *we* permit ourselves to use, but rather, we want to establish it because we have used premises that we are *in fact* permitted to use.

Here it will be useful to consider another potential pattern of attitudes towards modest and ambitious anti-sceptical projects. The pattern to consider here is that of *pessimism* about a modest anti-sceptical project, and *optimism* about an ambitious project. That is, maintaining that one can successfully establish that some contingent proposition is known, using only premises that would be permitted by the sceptic, but *not* permitted *by ourselves*. We have an immediate appreciation that someone holding *these* attitudes would be rather odd, to say the least. After all, it is not at all clear that the sceptic would permit a premise for the purposes of establishing that we have external

world knowledge, which *we* would not. Someone who held *that* pattern of attitudes would be very hard to understand at all. So much so, that it might be thought that there would be little to be gained by considering the possibility of someone who *did* hold these attitudes. To the contrary, I think that this strange possibility can help to illuminate my general concern with Pryor's divergent attitudes towards the prospects of success for ambitious and modest anti-sceptical projects. What makes this potential pattern of attitudes particularly useful to consider here is precisely *why* we are inclined to disregard it in the first place. What explains this? I think that a reasonably natural explanation might go as follows: there seems to be little reason for thinking that the sceptic could be satisfied that we can know, when *we ourselves* are not!

I think this situation reveals what lies behind the suspicion of Pryor's pattern of attitudes. In the case considered above, the strangeness consisted in being optimistic about *ambitious* responses to scepticism, and being pessimistic about modest ones. I think this strangeness can be explained by appreciating that, were the sceptic satisfied that we know some empirical proposition, then there would be good reasons for us to be satisfied too. On the assumption that the sceptic will *only* be satisfied that it has been established if *in fact* it *has* been, then *we* should be satisfied that it has too. And with that explanation in place I think we are now in a position to give a better general statement of our initial potential concern about Pryor's optimism and pessimism concerning modest and ambitious anti-sceptical projects, respectively. If the sceptic refuses to permit the use of some premises for the purposes of establishing that we know, or have justification for believing empirical propositions, then either they are *correct* to do so or not. If they are correct to do so, it is hard to see what use a modest anti-sceptical project could be in terms of solving the sceptical problem. There is little to be gained by pointing out that we can derive the falsity of sceptical conclusions from certain premises when the *use* of these premises for this purpose is correctly deemed to be inappropriate. But if, on the other hand, the sceptic's refusal to permit these considerations is *not* correct, then we should be able to show that this is the case. But this would not be shown by pointing out that we would permit the use of them, we would still need something more than what the modest anti-sceptical project provides.

For these reasons, it is not clear that distinguishing between what we may avail ourselves of in responding to scepticism will be helpful in understanding how the sceptical problem can be solved. In the remainder of this chapter and thesis, I set aside the distinction between ambitious and modest responses to scepticism, and move on to explore two further characterisations of approaches to solving the sceptical problem.

In contrast to the former distinction, both of these distinctions concern the aim of an response to scepticism, rather than the means by which this aim is supposed to be attained. I argue that when considered together, both of these provide a more useful dimension along which to assess whether a putative solution to the sceptical problem is successful.

1.2.3 Schiffer's Happy and Unhappy Face Solutions.

It is widely accepted that sceptical arguments present us with classical paradoxes. Schiffer has suggested that classical paradoxes, such as those which sceptical arguments present, may admit of a what he terms a 'happy-face' solution. Here is the explanation of what this sort of solution amounts to:

A happy-face solution to a paradox does two things, assuming that the propositions comprising the set really are mutually incompatible: first, it identifies the odd-guy-out, the member of the set that's not true; and second, it shows us why this spurious proposition deceived us, strips from it its patina of truth, so that we're not taken in by it again. (Schiffer, 2004, 178-9)

The point here is that there are two general adequacy constraints on providing a happy-face solution to sceptical problems. These can be thought of as two necessary steps for resolving the tension between the mutual inconsistency and apparent individual plausibility of the premises of sceptical arguments, and our claims to know various empirical propositions. These constraints, and question of whether they can be met, forms the crux of this thesis.

The first constraint is to identify a false claim, and thereby address the inconsistency partly constituting sceptical paradoxes. The second constraint is to provide an *error-theory*, or explanation for why the false claim was judged to be plausible. A happy-face solution to sceptical problems must meet both of these constraints. As Schiffer explains, it is necessary for a happy-face solution to

locate the false proposition in the set of mutually inconsistent propositions that we get by combining... [a sceptical argument's] premises with the denial of its conclusion, and the error theory is needed to explain why the sentence expressing the false proposition... deceptively appears to be stating a true proposition. (Schiffer, 1996, 325)

Schiffer maintains that some paradoxes do, in fact, have happy-face solutions. He points out that

[t]he paradox of the barber who shaves all and only those who don't shave themselves has a happy-face solution, for all sense of paradox disappears once we see that the existence of such a barber is logically impossible. (Schiffer, 2004, 179)

Schiffer has suggested that only happy-face solutions will be considered 'fully satisfactory'.⁴ When a 'happy-face' solution to a paradox is not possible, Schiffer claims, it is because of a 'glitch' in the concept involved in the paradox. These 'glitches', and the subsequent paradoxes, occur, he explains, when

[a]spects of the concept's underived conceptual role—the conceptual role the concept has regardless of whatever propositional attitudes one happens to have—are in tension, pull us in different directions, and there is nothing else in the concept or elsewhere to resolve that tension for us. (Schiffer, 2004, 179)

Schiffer suggests that the philosophical 'problem of free will' arises from a conceptual glitch, and consequently has no happy-face solution. This is, he says, due to

[o]ne aspect of the concept inclines us to apply the concept to certain paradigm cases, acts that are free if any are, whereas another aspect of the concept disinclines us to apply the concept to an act when we learn that the actor was caused to do what she did by factors over which she had no control. Further, there is nothing in the concept or elsewhere—no conceptual court of appeals—to resolve the tension by pronouncing one inclination legitimate, the other illegitimate. (Schiffer, 2004, 180)

Paradoxes that lack happy-face solutions have only what he terms 'unhappy-face' solutions. Since only happy-face solutions are fully satisfactory, an unhappy-face solution to a paradox will not be a fully satisfactory one. These less than fully satisfactory

⁴ See (Schiffer, 1996, 318). He suggests this, however, with a degree of 'coyness'. This, it would appear, is due to the fact that he does not think that sceptical paradoxes *can* have happy-face solutions. Consequently, he seems to want to keep open the possibility that other solutions can be fully satisfactory. I am unconvinced that solutions other than happy-face ones can be considered, in good faith, to be fully satisfactory. I explicate the reasons for this below.

solutions come in two sorts according to Schiffer. The first of these is what he terms a ‘weakly’ unhappy-face solution. Schiffer explains:

A weak unhappy-face solution is a mildly unhappy-face solution and says that a glitch-free version of the concept is possible which does the work we expected from the problematic concept[.] (Schiffer, 2004, 181)

What would a *weakly* unhappy-face solution to the problem of sceptical arguments such as DREAMING and ENVATTED BRAIN do? It would maintain two things. Firstly, that our inconsistent set of judgements arises from a genuine inherent incoherence, or ‘glitch’ in the concept of knowledge involved. Secondly, it would maintain that a ‘surrogate’ concept was available to replace knowledge, which would do the work we wanted from knowledge but without generating a paradox. It is worth considering an example of a paradox that Schiffer thinks does have a weakly unhappy-face solution. One would be the semantic paradoxes, as the following comments reveal:

As Tarski recognized, the semantic paradoxes can have no happy-face solution owing to features of our ordinary concept of truth in conjunction with certain logical concepts. But work on truth such as Kripke’s suggests that paradox-free alternative accounts of truth are possible, which don’t generate the semantic paradoxes, thus showing that the semantic paradoxes have a weak unhappy-face solution. (Schiffer, 2004, 181)

For those paradoxes that admit of neither a happy-face, nor a weak unhappy-face solution, Schiffer thinks we can have only what he terms ‘strong’ unhappy-face solutions. He qualifies this in the following way:

[A] strong unhappy-face solution is a very unhappy-face solution and says that no such unparadoxical surrogate can be fashioned. (Schiffer, 2004, 181)

When a paradox has only a strong unhappy-face solution it is because it involves a concept which is, in an important sense, inherently incoherent. An example of a paradox that can only have a strong unhappy-face solution, according to Schiffer, is the sorites paradox.⁵ It is worth noting that scepticism represents a *strong unhappy-face* solution to sceptical paradoxes. That is, our concept of knowledge, or justification is inherently incoherent. Schiffer is not a sceptic, however, and states that he thinks that sceptical paradoxes can have weakly unhappy-face solutions.⁶ I will consider, in the following

⁵ See (Schiffer, 2000).

⁶ See (Schiffer, 2004, 181 ff.).

sections, some reasons to think that anyone who does not wish to endorse scepticism should not be satisfied with a weakly unhappy-face solution. For my purposes, here, I will simply note that all things considered, a *happy-face* solution to sceptical paradoxes is preferable to an *unhappy-face* solution, for anyone who wishes to avoid endorsing scepticism. Crucially, then, Schiffer's distinction reveals a significant point about how to solve the sceptical problem. Ideally, solving the sceptical problem requires an *error-theory* that adequately explains why we might erroneously think that certain premises of sceptical arguments are plausible.

This concludes the discussion of Schiffer's distinction between happy-, and weakly or strongly unhappy-face solutions to sceptical problems. This distinction usefully distinguishes between three sorts of solution to the sceptical problem, and identifies the constraints that a response to scepticism must meet in order to constitute one. A happy-face solution would be, if attainable, obviously preferable to the alternatives. There are two constraints on such a solution. First, it must identify either (at least) one premise, or else the conclusion of sceptical arguments to be false, Second, it must provide an error-theory that explains away the false claim's apparent plausibility. In the next section, I outline and briefly discuss another distinction that has been proposed concerning solutions to the problem of sceptical arguments. This is the distinction between *theoretical* and *therapeutic* diagnoses of sceptical arguments.

1.2.4 Williams's Anti-sceptical Strategies

This brings us to the distinction concerning approaches to solving the sceptical problems that has been proposed by Williams (1991). Later on, I suggest that the considerations involved in this distinction have implications for how Schiffer's constraints on a 'fully satisfactory' happy-face solution to the sceptical paradoxes can be met. First, I outline the distinction between what Williams has termed 'theoretical' and 'therapeutic' diagnoses of sceptical problems, and then, subsequently, I describe what he terms the 'epistemologist's dilemma' confronting the latter.

Williams (1988; 1991; 2001) distinguishes between two approaches to solving the sceptical problem. A solution will involve either a *theoretical*, or else *therapeutic* diagnosis of the problem. In describing the former, theoretical, solutions, Williams explains:

I shall use '*theoretical diagnosis*' to refer to the strategy of attempting to uncover the sceptic's essential *epistemological* presuppositions. I shall never accuse the sceptic of incoherence. I shall not argue that his problems are

pseudo-problems. On the contrary, I think they are fully genuine, *but only given certain theoretical ideas about knowledge and justification*. (Williams, 1991, 37)

The aim of a theoretical solution to sceptical problems is to “tell us what theoretical commitments the problem of the external world, and its attendant skepticism, are grounded in” (Brady and Pritchard, 2005, 355). The anti-sceptical strategy involved in theoretical solutions to the sceptical problem is to “show that sceptical arguments depend essentially on theoretical commitments that are *not* forced on us by our ordinary ways of thinking about knowledge, justification, and truth” (Williams, 1991, 31-2). An approach to solving the sceptical problem by means of a theoretical diagnosis will maintain that the epistemic principles which seem to underwrite the premises of sceptical arguments

are unnatural ones—which is to say, they are not forced on us by considerations that any reasonable person would have to accept as compelling. Instead, they only arise if one has already taken on board a variety of controversial theoretical commitments...[i.e.] a set of—on examination, highly dubious—philosophical theories. If we reject those presuppositions, we can with a good conscience simply reject the [sceptical argument.] (Rudd, 2008, 305)

Here, Williams regards it as a major advantage of a theoretical diagnosis of sceptical problems, that it circumvents any concerns we might have about whether we can demonstrate that a sceptical argument is not sound, or, *refute* it. He explains this in the following way:

[I]f we hand the sceptic his... presuppositions, there is no refuting him: this is what the conditional correctness of scepticism consists in. And if we miss the way seeming platitudes can be used to smuggle in the crucial epistemological ideas, the conditional correctness of scepticism will make scepticism seem inevitable. But if we keep these ideas clearly in view the situation is changed: there is no danger in conceding that the sceptic cannot be refuted on his own terms if those terms are not ones we are bound to accept. (Williams, 1988, 417)

Williams contrasts theoretical diagnoses of sceptical problems with what he terms ‘therapeutic’ diagnoses. Central to the distinction is Williams’ observation that “arguments that appeal to something deep in our nature contrast with arguments that turn on

special theoretical presuppositions” and that “to appeal to something deep in our nature, an argument for scepticism need not be completely presuppositionless. It must, however, exploit only the most deeply embedded features of our ordinary conception of knowledge” (Williams, 1991, 18).

Theoretical diagnoses maintain that sceptical arguments depend upon, or ‘turn on’ special theoretical presuppositions. They hold that the premises of these arguments are motivated by principles which are extraneous of our ordinary ways of thinking about knowledge or justification. A solution to the sceptical problem need not do so. Instead, a solution may hold that sceptical arguments are motivated only by ‘natural’ ways of thinking about knowledge. Solutions that do this involve what Williams terms therapeutic diagnoses. It is characteristic of them, he explains, to maintain that

sceptical arguments as highly ‘natural’ or ‘intuitive’ in the sense of minimally dependent on contentious theoretical ideas. Thus Stroud claims that ‘when we first encounter... sceptical reasoning... we find it immediately gripping’, which he takes to indicate that such reasoning ‘appeals to something deep in our nature’. (Williams, 1988, 416)

The distinction between theoretical and therapeutic solutions to sceptical problems consists in the fact that theoretical maintain “that the sceptic depends essentially on distinctive theoretical commitments not clearly implicit in our ordinary handling of epistemic concepts” (Williams, 1991, 32). Since therapeutic solutions, by definition, deny that sceptical arguments are based on contentious theoretical considerations, therapeutic solutions involve arguing that sceptical problems are

dissolved by showing that the sceptic doesn’t or can’t mean what he seems to mean, perhaps even that he does not succeed in meaning anything at all. It is not enough to show that the case for scepticism is less than compelling: he has to show that no coherent problem was ever presented. (Williams, 1991, 32)

Any therapeutic diagnosis, on Williams’ view, “treats sceptical problems as *pseudo-problems* generated by misuses or misunderstandings of language. On this approach, sceptical claims and arguments do not really make sense” (Williams, 2001, 146). The characteristic aim of a therapeutic solution to sceptical problems is that of “exposing epistemological problems as illusory, thus making them disappear without theoretical residue” (Williams, 2001, 253).

Williams maintains that any response to scepticism that is not a theoretical diagnosis of scepticism, is unable to avoid what he terms the ‘epistemologist’s dilemma’.⁷ On his view, there is a choice of unpalatable alternatives for any diagnosis of the problem according to which sceptical arguments are based on epistemic principles that are part of our ordinary ways of thinking about knowledge. This is because Williams thinks that any therapeutic diagnosis of the problem can avoid embracing scepticism only by either maintaining that we do not really understand scepticism, or else by giving up on the epistemic principle that we otherwise ordinarily accept. I think it will be instructive to distinguish these two sorts of therapeutic diagnoses. Let’s call those that maintain that we do not really understand sceptical arguments *confusion-therapeutic*, and those that involve abandoning a platitudinous epistemic principle *concessive-therapeutic*.

Adopting either a confusion-therapeutic or concessive-therapeutic diagnosis of scepticism represents, according to Williams, an unattractive option for an anti-sceptical epistemologist. To see why, let’s consider the latter option first. A concessive-therapeutic response to scepticism involves accepting that sceptical arguments are ‘natural’, in the sense outlined above, and avoiding accepting sceptical conclusions by abandoning those platitudes about knowledge on which it relies. Yet, Williams argues, it is hard to see how this amounts to anything other than *conceding* the very point of sceptical arguments. He explains the problem facing any concessive-therapeutic response in the following way:

Clearly, the sceptic triumphs if the platitudes he makes use of turn out to be impossible to deny. But even *if we do manage to deny them*, the sceptic can argue that to abandon [a platitude]... is to indicate our willingness to settle for something less than knowledge of the world, as we have always understood it; and *this is to concede that knowledge as we have always understood it* is indeed beyond us. So the sceptic triumphs either way. (Williams, 1988, 417, my emphasis.)

Here, Williams is describing what he takes to be the first horn of the ‘epistemologist’s dilemma’ facing therapeutic solutions to scepticism. If sceptical conclusions *do* follow from premises that are motivated by deeply held epistemic principles, and the sceptical conclusion is to be rejected, then these principles must be *extricated* from our concept of knowledge. This amounts to a concession of the sceptical position—it is to concede that we do not know any empirical propositions, *in the sense that we thought we did*. I think

⁷ For critiques of Williams’ optimism concerning theoretical solutions to sceptical problems, see (Hetherington, 1994), and (Rudd, 2008).

that this aspect of the dilemma facing therapeutic responses to sceptical arguments has important implications for happy-face resolutions of sceptical paradoxes. And so it is worth stressing the point. The following comments, for example, neatly captures the inherent problem in trying to find a happy-face solution in terms of a concessive-therapeutic diagnosis:

For we can no longer plausibly argue that the skeptic is enforcing some arcane, remarkable, or implausible requirement on perceptual knowledge. Instead, *one is faced with arguing against an epistemic principle which one, along with most, if not all us, adhere to in otherwise normal circumstances.* The dialectical advantage thus goes to the skeptic rather than us. (Pritchard and Ranalli, 2013, 353-4, my emphasis.)

To avoid the first horn of the dilemma—conceding previously endorsed epistemic principles, and, subsequently, precluding a happy-face solution— a non-theoretical solution will need to be confusion-therapeutic one. This involves attempting to avoid endorsing sceptical conclusions by maintaining that the use of platitudinous principles in sceptical arguments is somehow incoherent, or unintelligible. Even these approaches, Williams argues, face problems of their own. If we want a happy-face anti-sceptical solution to sceptical paradoxes that is not a theoretical diagnosis, then it seems we will need to maintain that sceptical arguments are ‘less than fully intelligible’. Williams explains what he takes to be the problem with such a confusion-therapeutic diagnosis in the following way:

The trouble is... [that sceptical arguments] do not seem defective in point of intelligibility, particularly not to those who find them strikingly intuitive. In fact, it is not clear that one could hold that such arguments invoke only platitudes and still find them less than fully intelligible. Would they not have to inherit their lack of intelligibility from the platitudes they make use of? So... the evident intelligibility of what the sceptic claims counts against any theory... that calls its intelligibility in question. (Williams, 1988, 421-2)

This gives us the second horn of the epistemologists dilemma facing any non-theoretical solution to sceptical problems—specifically, confusion-therapeutic ones. If sceptical conclusions *do* follow from premises that are motivated by deeply held epistemic principles, and the sceptical conclusion is to be rejected, then, it seems, we do not really

understand sceptical arguments at all. Yet, Williams argues, this just seems to be an implausible position to maintain. As he explains:

Our sense that we *do* understand the sceptic—well enough, for example, to understand how we might argue *against* him— will eventually wear down the credibility of theories that imply that we don't. (Williams, 1991, 18)

The epistemologist's dilemma, then, can be characterised in the following way:

[I]f scepticism is really based upon ordinary platitudes, then the claim that we really don't understand the sceptic, or a solution that requires a revision of an ordinary platitude, will, in the end, be a round about way of agreeing with the sceptic that, as Williams puts it, "knowledge of the world" as we ordinarily (though tacitly) understand it, is impossible. (Buchanan, 2002, 77)

The supposed upshot of the epistemologist's dilemma, then, is that *both* confusion-, and concessive-therapeutic diagnoses of sceptical arguments cannot represent happy-face resolutions of sceptical paradoxes. I agree that a concessive-therapeutic solution is unable to represent a happy-face resolution, for the reasons considered. In virtue of *conceding* the platitudinous epistemic principle, such a solution will necessarily be a weakly unhappy-face resolution of sceptical paradoxes.

I do not think, however, that a confusion-therapeutic diagnosis is unable to be part of a happy-face solution. This is because I do not think, as Williams' does, that maintaining that sceptical arguments are 'less than fully intelligible' implies that our ordinary ways of thinking about knowledge are too. If it did, then confusion-therapeutic diagnoses would, ultimately, amount to either weakly or strongly unhappy-face resolutions. But it need not. This is, then, the second crux of this chapter and thesis: that is possible to both maintain that sceptical arguments rely on platitudinous epistemic principles, and to explain the source of our error (or *confusion*) with respect to sceptical arguments in other terms. And I will return to explore these point in depth in the second half of this thesis. In chapters 4 and 5, for example, I investigate the whether an anti-sceptical error theory for sceptical paradoxes can be given in *semantic* terms. In chapter 6, I consider whether a happy-face resolution of sceptical paradoxes can be had by means of an error theory given in *cognitive* terms.

That concludes the discussion of both Williams' distinction between theoretical and therapeutic solutions to sceptical problems, and what he terms the epistemologist's

dilemma confronting the latter. In the next section, set out some ways in which I think the distinctions drawn by Schiffer and Williams can be brought together. I suggest that this reveals some significant implications for the prospects of providing an anti-sceptical happy-face resolution of sceptical paradoxes.

1.3 Theoretical and Therapeutic responses: Happy-face Solutions?

In the previous section, I argued that confusion-therapeutic diagnoses of sceptical arguments *are not* inherently incapable of providing happy-face resolutions of sceptical paradoxes. Now, in this section, I aim to pave the way for the ensuing chapters of the thesis by suggesting that theoretical diagnoses face significant obstacles to providing happy-face resolutions.

1.3.1 Second order constraints on theoretical responses.

The primary crux of this chapter and thesis is that in order to be a happy-face solution to sceptical paradoxes, and subsequently, a fully satisfactory solution to the sceptical problem, a response to scepticism must fulfill two conditions. First, it must identify the false claim within the apparently inconsistent and plausible claims presented by sceptical arguments. Second, it must provide an error-theory, i.e. an account of why it is that we might mistakenly judge this false claim to be plausible. Theoretical anti-sceptical solutions will, in general, attempt to meet the first condition by identifying either of a sceptical argument's premises as false. Subsequently, they will attempt to meet the second condition, in general, by giving an account of an epistemic principle underwriting the premise that is identified as false.

I think, however, that the considerations pertaining to concessive-therapeutic solutions involved in the epistemologist's dilemma, provide good reasons to think that a theoretical solution will not obviously provide happy-face solutions. Recall that concessive-therapeutic solutions are revisionary, in the sense of replacing one way of thinking about knowledge in favour of another, and for which paradoxes will not arise. As such, these are clearly weakly unhappy solutions. But can a theoretical response to scepticism constitute a happy-face solution to the sceptical problem? It can only if it is able to adequately explain why it is that we get confused by sceptical arguments in terms of a principle that is *not* a part of ordinary way of thinking about knowledge. If no explanation of this

1.3. THEORETICAL AND THERAPEUTIC RESPONSES: HAPPY-FACE SOLUTIONS?

sort is forthcoming from a theoretical anti-sceptical solution, then it will not obviously provide a happy-face solution to sceptical paradoxes. Yet, given what I take to be a plausible assumption about what providing this explanation *would* involve, I suggest that theoretical responses face a significant obstacle to constituting a fully satisfactory solution to the sceptical problem.

How might a theoretical solution explain why we were misled, by a contentious epistemic principle, into erroneously finding a sceptical argument's false premise to be plausible? I assume that an explanation of this sort will involve making one of at least two sorts of claim:

TE1 We mistakenly thought that the principle *was* platitudinous, i.e. a principle that we ordinarily adhere to in our ordinary ways of thinking about knowledge.

TE2 The principle, whilst *not* actually a platitude, is *sufficiently similar* to one that is, such that we mistakenly thought sceptical arguments relied on the latter rather than the former.

A theoretical solution might constitute a happy-face solution by maintaining either that we thought an epistemic principle was a platitude, when it was not, or that we thought that an actual epistemic platitude motivated sceptical arguments, when it did not. I suggest, however, that there are certain second-order considerations which cast doubt on the potential to provide a theoretical happy-face resolution of sceptical paradoxes in terms of either TE1 or TE2. Can an explanation, given in terms of either of these claims, be distinguished, in principle, from a concessive-therapeutic solution? If not, then I think there will remain a second-order obstacle to thinking that a theoretical anti-sceptical project constitutes a happy-face solution. This is because, as noted above, concessive-therapeutic solutions involve revising our ordinary ways of thinking about knowledge, which precludes their constituting fully-satisfactory happy-face solutions.

So, is there a way to distinguish, in principle, a theoretical *explanation* for our propensity to erroneously find sceptical premises plausible, from a concessive-therapeutic *diagnosis* of sceptical arguments? I think that there *would* be only if there were good reasons to think, independently of sceptical paradoxes, that the epistemic principle underwriting sceptical arguments *never really was* a platitude. Yet, if we had good reasons to think *that*, then presumably we would already have available facts which showed, independently of scepticism, that these principles were *appreciably not* platitudinous. But, in this case, any attempt to explain why we got confused by sceptical arguments, in terms of claims such as TE1 or TE2, will be significantly undermined. This is because it

will remain unclear *why* we would either *think it was a platitude*, or *confuse it with a principle that is*, if the epistemic principle underwriting sceptical arguments is *appreciably* not platitudinous independently of consideration of sceptical arguments. In either case, since our confusion when considering sceptical argument remains unexplained by such a theoretical diagnosis, it will be unclear whether it provides a happy-face resolution of sceptical paradoxes.

In order to be considered a happy-face solution, a theoretical solution to scepticism must explain our error with respect to sceptical paradoxes, and distinguish this solution from *in fact* providing only a therapeutic solution. I have suggested, however, that a theoretical—in other words, *epistemological*—error theory might only be distinguishable from a concessive-therapeutic solution in a way which could undermine the very credibility of error theory itself. In this way, I think that theoretical diagnoses of sceptical arguments face significant obstacles constituting fully satisfactory, happy-face resolutions to sceptical paradoxes.

In the remainder of this chapter, I focus on applying the preceding considerations concerning theoretical diagnoses to Pryor's own attempt to provide a solution to sceptical problems. In sections 1.3.2 and 1.3.3, I examine the details of Pryor's proposed solution to scepticism in terms of the two conditions on happy-face resolutions. Finally, in section 1.3.4, I suggest—in line with the foregoing reasoning—that Pryor's proposal does not obviously constitute a happy-face resolution to sceptical paradoxes.

1.3.2 Pryor's anti-sceptical project.

I turn now to explicate how Pryor's own proposed anti-sceptical project attempts to satisfy the first condition on happy-face solutions. This is the requirement that the solution identifies the false claim within the set of claims comprising the paradox. Here, I set out some features of Pryor's proposal that are relevant to meeting the second condition on a happy-face solution with respect to DREAMING and ENVATTED BRAIN: an explanation for the apparent plausibility of the false claim.

Which of the premises of sceptical arguments, such as DREAMING, or ENVATTED BRAIN might Pryor maintain is false? Pryor, I take it, maintains that it is premises (4) and (7), respectively, that are false. To see how Pryor might could account for the falsity of these premises, it will be instructive to first look at how he maintains that we can have justification for believing, and potentially know, some empirical proposition.

Pryor distinguishes between two sorts of justification to believe a proposition, which he terms 'mediate' and 'immediate' justification. These he explains these in the following

1.3. THEORETICAL AND THERAPEUTIC RESPONSES: HAPPY-FACE SOLUTIONS?

way:

Say that you are “mediately justified” in believing p iff you’re justified in believing p , and this justification rests in part on the justification you have for believing other supporting propositions. Say that you are “immediately justified” in believing p , on the other hand, iff you’re justified in believing p , and this justification doesn’t rest on any evidence or justification you have for believing other propositions. (Pryor, 2000, 532)

Pryor’s anti-sceptical project establishes that we can justifiably believe and know various empirical propositions in terms of an account of ‘immediate’ perceptual justification. This account holds that

whenever you have an experience as of p , you thereby have immediate *prima facie* justification for believing p ... Your experiences do not, in the same way, give you immediate *prima facie* justification for believing that you are dreaming, or being deceived by an evil demon, or that any of the skeptic’s other hypotheses obtain. (Pryor, 2000, 536)

On Pryor’s view, if S has an experience ‘as of p ’, for some empirical proposition p , then S has a sort of perceptual justification to believe that p , and can thereby count as knowing that p . Pryor admits that, in presenting this account of perceptual justification, he does not attempt to explain in detail

how we’d go about ruling out a skeptical hypothesis, either in the case where we have positive evidence that it obtains, or in the ordinary case, where we have no evidence in favor of it but we’re just curious whether or not it obtains. But a few brief remarks may help clarify my position. (Pryor, 2000, 546)

My purpose here is to assess whether Pryor’s anti-sceptical proposal, understood as a theoretical diagnosis, constitutes a happy-face solution to sceptical paradoxes, such as ENVATTED BRAIN and DREAMING. To do so, it identifies premises (4) and (7), for example, as the false claims partly comprising the paradoxes. It is important, then, to understand what sort of explanation Pryor could give for the falsity of these premises. How might we, on Pryor’s view, have justification to believe, or know, that sceptical hypotheses are false? In order to see how, it will be instructive to consider the following comments, in which Pryor frames his answer to that question:

Suppose U is some hypothesis such that the only reasons you have for believing U to be false presuppose the truth of p , but if you were to acquire evidence for U , that would defeat or undermine your justification for believing p . (For example, let p be some body of perceptual beliefs, and let U be a belief like *I am a brain in a vat* or *My senses are unreliable*. (Pryor, 2000, 546)

Pryor's key consideration here is what he takes to be the pertinent relation between our 'immediately' *prima facie* justified beliefs concerning various empirical propositions, and our beliefs that sceptical hypotheses are false. Namely, that the only reasons we have for believing that sceptical hypotheses, such as the *biv* or *dreaming* hypotheses, are false 'presuppose' some empirical propositions. How, on his account, do our beliefs with respect to sceptical hypotheses then amount to knowledge?

Before considering the account Pryor gives of the false claim involved in sceptical arguments such as ENVATTED BRAIN, or DREAMING, it will be instructive to consider what he means by the term 'bad' with respect to sceptical hypotheses. His explanation of this term begins as follows:

Say that an alternative to p is a "bad" alternative just in case it has the special features that characterize the skeptic's scenarios—whatever those features turn out to be. Different skeptical arguments will rely on different accounts of what makes an alternative "bad". (Pryor, 2000, 527)

There are, Pryor thinks, two main accounts one might give of the 'special features' characterising sceptical hypotheses. He explains the first of these as follows;

Say that some grounds E you have "allow" a possibility q iff the following counterfactual is true: if q obtained, you would still possess the same grounds E So we might want to count a hypothesis as "bad" for the purpose of a skeptical argument just in case it is (and is recognized to be) incompatible with what you purport to know, but it is nonetheless "allowed" by your grounds E , in the sense I described. (Pryor, 2000, 527)

The *biv* hypothesis, for example, could be characterised as 'bad' in this way. It is incompatible with many empirical propositions that one might purport to know (that *you have hands*, for example). On the assumption that your perceptual experiences comprise your grounds E for believing that you have hands, it seems that if the *biv* hypothesis were true with respect to you, then you would still possess the same grounds E . So, *biv* is

1.3. THEORETICAL AND THERAPEUTIC RESPONSES: HAPPY-FACE SOLUTIONS?

‘allowed’ by your grounds, and ‘bad’ with respect to any proposition that is incompatible with it, for example, that *you have hands*.

The *dreaming* hypothesis, however, is not ‘bad’ in *this* way with respect to practically any proposition you might purport to know. This is because very few propositions are incompatible with it being true. Pryor is aware of this, and explains the second account of a ‘bad’ possibility as follows:

[It] does however introduce a non-standard explanation of your experiences. And this explanation would *undermine* the support your experiences give you for your perceptual beliefs—in the sense that, if you were to learn that you are dreaming, then you would have reason to doubt that your experiences were a trustworthy basis for beliefs about the external world. So we might want to count a hypothesis as “bad” for the purposes of a skeptical argument if it could undermine your experiences, in this sense. (Pryor, 2000, 527)

Which of the three claims, then, of comprising sceptical paradoxes would Pryor’s anti-sceptical project identify as false? Above it was noted that, with respect to ENVATTED BRAIN and DREAMING, Pryor seems to be inclined to regard the first premise of each argument to be false. So what could explain, on Pryor’s account, our knowledge that we are neither brains in vats, nor currently dreaming? On his view, I take it, where U refers to a sceptical hypothesis, these premises are false for the following reasons:

[Y]our justification for believing *p* does give you justification for believing not-U. However, because U is a potential defeater or underminer of your justification for believing *p*, any evidence you acquired in favor of U would defeat (or at least contribute towards the defeat of) your justification for believing *p*. (Pryor, 2000, 546)

If someone has *immediate prima facie* justification to believe that *p*, then they thereby have *mediate* justification to believe that they are not dreaming, or a brain in a vat, so long as they do not acquire evidence in favour of either these hypotheses. On Pryor’s view, since almost every subject has immediate *prima facie* for believing some empirical proposition, they also have *mediate* justification for believing both that the *brain in a vat*, and that *dreaming* hypothesis is false. Consequently, almost every subject can be said to know that, with respect to themselves, these hypotheses are false. And this, on Pryor’s view, is what allows us to identify both the sceptical premises, (4) and (7), as false.

1.3.3 Sceptical Principle.

In this section, I consider in detail Pryor's account of the epistemic principle that motivates the false sceptical premises (4) and (7), which deny that we know that sceptical hypotheses are false. Above I suggested that a theoretical anti-sceptical project must identify the epistemic principle motivating sceptical premises in order to be a happy-face resolution of sceptical paradoxes. It will be instructive then, for the purposes of assessing whether Pryor's response to scepticism could be a happy-face one, to examine how he identifies this principle. Pryor's approach to identifying the principle which motivates these problematic premises involves undertaking a reconstruction of a typical sceptical argument. The first premise that Pryor attributes to the sceptic in his reconstruction of their reasoning is the following:

(P1) Either you don't know you're not being deceived by an evil demon; or, if you do know you're not being deceived, it's because that knowledge rests in part on things you know by perception.⁸

In this reconstruction of the sceptical argument, Pryor uses the hypothesis that one is being deceived by an 'evil' demon.⁹ This hypothesis has the same function, in Pryor's reconstructed sceptical argument, as the *biv* and *dreaming* hypotheses have in the arguments ENVATTED BRAIN and DREAMING, respectively. The sceptic might argue, along the lines of (P1), that either S knows that they are not a brain in a vat; or, if they do, then it is because that knowledge rests in part on things S knows by means of perception.

The next premise that Pryor attributes to the sceptic in his reconstruction of their argument is the following:

SPK If you're to know a proposition p on the basis of certain experiences or grounds E , then for every q which is "bad" relative to E and p , you have to be in a position to know q to be false in a non-question-begging way—i.e., you have to be in a position to know q to be false antecedently to knowing p on the basis of E .¹⁰

For Pryor, SPK encapsulates a significant sceptical idea. This is, according to him, that "in order for us to know anything on the basis of perception, we *first* have to know we're not being deceived" (Pryor, 2000, 524). Pryor is careful, however, to clarify what

⁸ (Pryor, 2000, 527).

⁹ See (Descartes, 1982).

¹⁰ (Pryor, 2000, 528).

1.3. THEORETICAL AND THERAPEUTIC RESPONSES: HAPPY-FACE SOLUTIONS?

he means by both ‘antecedently’ and ‘first’ here; the sense of priority he has in mind is *epistemic*. He goes on to explain this as follows:

Your justification for believing p_1 is antecedent to your justification for believing p_2 just in case your reasons for believing p_1 do not *presuppose or rest on* your reasons for believing p_2 . Your reasons for believing p_1 can not *beg the question* whether p_2 .¹¹

Next, Pryor attributes the following premise to the sceptic:

(P2) The hypothesis that you’re being deceived by an evil demon is “bad” relative to *any* course of experience E and perceptual belief p . (Pryor, 2000, 528).

This premise encapsulates the apparently uncontroversial idea that sceptical hypotheses, such as the *brain in a vat*, and *dreaming* hypotheses, are ‘bad’, in the sense discussed above, with respect to all our perceptually based beliefs in empirical propositions. In his reconstruction of the sceptical argument, Pryor puts (P2) and SPK together to put forward the following premise:

(P3) If you’re ever to know anything about the external world on the basis of your perceptual experiences, then you have to be in a position to antecedently know you’re not then being deceived by an evil demon.¹²

With these premises in place, Pryor argues, the sceptic could derive their premise that we do not know that sceptical hypotheses are false, by *reductio*. Suppose, for example, that S knows that a sceptical hypothesis, sh , is false. The considerations involved in (P1) imply that S’s knowledge that not- sh rests in part on knowledge they have by means of perception. Call this piece of perceptual knowledge p^* , as Pryor does. The considerations involved in (P3) imply that S knows that p^* only if S is in a position to ‘antecedently’ know that not- sh ; or, S’s reasons for believing that p^* do not ‘presuppose or rest on’ their reasons for believing that not- sh .

Since p^* is a piece of *perceptual* knowledge—by (P1) and the assumption that S knows that not- sh ,—it follows from SPK and the considerations involved in (P2), that S’s reasons for believing that not- sh do *not* presuppose or rest on S’s reasons for believing that p^* . So, the assumption that S knows that not- sh implies that this piece of knowledge both

¹¹ (Pryor, 2000, 525).

¹² (Pryor, 2000, 528).

does and *does not* rest in part on things S knows by means of perception. On this basis, Pryor maintains, we can infer that the assumption that S knows that not-*sh* is false, and arrive at the following sceptical premise:

(P4) You do not know that you are not being deceived.¹³

This completes Pryor's reconstruction of the considerations involved in the sceptic's reasoning. On this account, the sceptical premise (P4) is motivated by the epistemic principle SPK. This sceptical premise differs only from (4) and (7) in that it involves the hypothesis that one is being deceived. This hypothesis is functionally equivalent to the sceptical hypotheses involved in both ENVATTED BRAIN and DREAMING. Consequently, we can infer that the apparent plausibility of (4) and (7) is explained, on Pryor's view, by the epistemic principle SPK. He is explicit, for example, in claiming that "if you look at informal presentations of the skeptic's reasoning, you'll find that these *do* often rely on some principle like SPK" (Pryor, 2000, 529). In the next section, I suggest that Pryor's proposed solution, understood as a theoretical diagnosis of sceptical arguments, does not obviously provide us with a happy-face resolution of sceptical paradoxes.

1.3.4 The Inadequacy of Pryor's Project.

Pryor's solution attempts to meet the first condition on a happy-face resolution by identifying the minor premises (4) and (7) of sceptical arguments DREAMING and ENVATTED BRAIN, respectively, as false. It does so in terms of his account of our having immediate perceptual justification for believing that various empirical propositions, and, subsequently, our having mediate justification for believing that sceptical hypotheses are false. His solution attempts to meet the second condition, by giving an account of the apparent plausibility of minor sceptical premises in terms the epistemic principle SPK.

Pryor's response can plausibly be interpreted as an attempt to provide a theoretical diagnosis of sceptical arguments. It would only be therapeutic *if* it held that SPK is a principle that we adhere to in our ordinary ways of thinking about knowledge. And I do not think that is what Pryor intends to. But if Pryor's theoretical solution is to be a happy-face one it will need to have the following two features. First, it will maintain that the epistemic principle SPK is *not* one that we adhere to in our ordinary ways of thinking about knowledge. Second, it will explain why SPK misled us into erroneously finding it plausible to deny that we know sceptical hypotheses are false.

¹³ cf. (Pryor, 2000, 528).

1.3. THEORETICAL AND THERAPEUTIC RESPONSES: HAPPY-FACE SOLUTIONS?

Can Pryor adequately explain why SPK misleads us, in this way, if this principle is not a part of ordinary way of thinking about knowledge? I argue that he cannot, given the only plausible assumption that this will involve making one of either two claims corresponding to TE1 or TE2. The first is that, we mistakenly thought that we ordinarily do adhere to SPK in our ordinary ways of thinking about knowledge. The second is that, whilst we do *not actually* adhere to SPK in our ordinary ways of thinking about knowledge, it is sufficiently similar to a principle that *we do*, and we mistook SPK for this latter principle. Suppose that Pryor's solution puts forward an explanation in terms of either one of these claims. Can this explanation be distinguished, in principle, from *in fact* rejecting the ordinarily adhered to SPK, and, subsequently, *revising* our concept of knowledge? If not, then it will remain an open question whether Pryor's solution constitutes a happy-, or weakly unhappy-face solution to sceptical problems.

Is there a way to distinguish an explanation for why SPK misled us with respect to sceptical arguments, in terms such as TE1 or TE2, from actually abandoning the otherwise platitudinous SPK? I suggest that there is only if we have good reasons to think, independently of sceptical paradoxes, that SPK *never really was* adhered to in our ordinary ways of thinking about knowledge. In other words, Pryor's theoretical solution can be distinguished from a concessive-therapeutic—and subsequently, a weakly unhappy-face one—only if we can avail ourselves of facts which showed, independently of considerations of scepticism, that SPK *was not* platitudinous to begin with. Yet, if SPK were appreciably *not platitudinous*, independently of sceptical arguments, then Pryor's solutions potential explanation for our confusion in terms such as TE1 or TE2 would be undermined. Such explanations would either involve claiming that we mistook SPK for a platitude, or that confused it with one that is a platitude. And neither explanation, however, has much plausibility when it is maintained that SPK is appreciably not platitudinous independently of sceptical arguments. In each case it remains unexplained why we would either *think that* SPK *was a platitude*, or *confuse* SPK *with another principle that is*. Consequently, the explanation provided by Pryor's theoretical anti-sceptical project for why we erroneously judge sceptical premises such as (4) and (7) to be plausible will be significantly undermined.

I have suggested that Pryor's theoretical solution must do two things in order to be considered a happy-face resolution of sceptical paradoxes. It must explain the erroneous plausibility of denying that we know sceptical hypotheses are false, in terms of the non-platitudinous epistemic principle SPK. And it must distinguish this from a rejection of the otherwise platitudinous SPK. Yet, I have suggested, Pryor's solution is seemingly

only able meet the latter condition in a way that will undermine its attempt to meet the former condition. I conclude that there are serious obstacles to thinking that his theoretical response to scepticism will provide us with a happy-face resolution of sceptical paradoxes.

1.4 Concluding remarks.

It remains here to draw together the preceding discussion to appraise the situation concerning the sceptical problem. The problem of scepticism is that sceptical arguments represent paradoxes, and, subsequently, threaten to reveal that our ways of thinking about knowledge, or justification are incoherent. A satisfactory solution to this problem is required to do two things: identify the false claim amongst a sceptical argument's premises and the negation of its conclusion, and *adequately* explain why this false claim was judged to be plausible.

What can be said about the prospects of achieving this? In this chapter, I have argued that various attempts to distinguish approaches to solving the problem do not provide a way to satisfactorily solve the sceptical problem. In particular, it was argued both that concessive-therapeutic responses are unable to constitute happy-face solutions *in principle*, and that theoretical responses face significant second-order obstacles to providing an adequate error-theory. Is there a way to overcome these obstacles, and provide an *adequate* error-theory concerning sceptical paradoxes?

The rest of this thesis is concerned with providing a framework for providing a positive answer to this question by drawing upon the points raised in this chapter. Specifically, that an adequate error-theory concerning sceptical paradoxes will need to be given in terms that are not purely *epistemological*. Instead, it will be argued, that a fully satisfactory, happy-face resolution of sceptical paradoxes can be had when an error-theory concerning sceptical paradoxes is given in *cognitive* terms.

The thesis subsequently breaks in two subsequent parts. The first of these, comprised of chapters 2 and 3, consists in providing diagnostic discussion of sceptical arguments in terms of sceptical hypotheses, and the epistemic principles underwriting them, respectively. The second part, comprised of chapters 4, 5 and 6 consists in assessing the potential to find happy-face resolutions of sceptical paradoxes in rival epistemological views of contextualism, and non-sceptical invariantism.

SCEPTICAL HYPOTHESES.

2.1 Introduction.

In the previous chapter, I outlined the central issue with which I am concerned in this thesis: the problem of sceptical paradoxes. This is the problem that arguments can be put forward to challenge our putative knowledge of any external world proposition at all. These arguments result in paradoxes; they are apparently sound, but their conclusions are unacceptable. My aim in this thesis is to explore the potential for providing a ‘happy-face’ solution to sceptical paradoxes. I argued, in chapter 1, that the most promising route to a happy-face solution lies in a *therapeutic*, rather than a theoretical diagnosis, of how sceptical arguments present us with paradoxes. As we saw, in order to be happy-face resolution, a happy-face resolution of sceptical paradoxes must provide an error-theory that accounts for our confusion regarding sceptical arguments. I set out, in chapter 6, my proposal for an anti-sceptical error-theory for sceptical paradoxes that is compatible with a therapeutic approach. But before either setting out this error-theory, or critically assessing existing proposals (see chapters 4 and 5), I think it will be instructive to delve deeper, and provide some diagnoses of what sceptical arguments involve. To this end, in both this chapter and chapter 3, I explore further the diagnostic approach to sceptical arguments from ignorance, by analysing them in terms of two elements: sceptical hypotheses, and epistemic principles.

In this chapter, I am concerned specifically with an attempt to understand sceptical hypotheses. My discussion addresses the following question: what are the conditions

that a hypothesis must satisfy in order to be an *effective sceptical hypothesis*? I begin, in section 2.2 by looking at an account of sceptical hypotheses that has informed the debate proposed by DeRose (1995), which attempts to account for sceptical hypotheses in terms of a supposed sensitivity condition on knowledge. After outlining the sensitivity account of sceptical hypotheses, I move on, in section 2.3, to outline a proposed account of necessary conditions on sceptical hypotheses put forward by Beebe (2010): the experiential constraints. In section 2.4, I introduce a novel case, and argue that it represents a counter-example to the sensitivity account, and experiential constraints. In section 2.5, I outline and argue against the necessity of a further condition proposed by Beebe: the indication constraint. In sections 2.6 and 2.7, I draw upon the considerations involved in the indication constraint, and related accounts of sceptical hypotheses put forward by Pryor (2000) and Williamson (2000), to develop a positive account of a necessary condition on effective sceptical hypotheses in terms of a subject's basis for belief.

2.2 Sceptical Hypotheses and Sensitivity.

Before discussing DeRose's proposal, it is important to clarify a couple of aspects of what is meant by an 'effective sceptical hypothesis'. On the one hand, a hypothesis is sceptical if it can be used to challenge our epistemic position with respect a whole range of propositions. In this sense, the hypothesis that *my bicycle has just been stolen* is *not* a sceptical hypothesis. Yet, as we saw in chapter 1, a hypothesis need not call into question *every* empirical proposition in order to count as sceptical. The hypothesis that *the world came into existence five minutes ago*, for example, does not challenge my putative knowledge of having hands. It does, however, seem to be a sceptical hypothesis insofar as it calls into question my epistemic position with respect to a range of propositions (for example, any proposition about the existence of dinosaurs). In this way, being a sceptical hypothesis would appear to be a matter of degree, and some, like those involving considerations of brains in vats, are maximally sceptical while others might be just challenges a specific range of propositions.

The second important aspect of what is meant by an effective sceptical hypothesis concerns our epistemic position with respect to *it*. An effective sceptical hypothesis is one that can be used to challenge a range of propositions *effectively*. To explain this, it is useful to consider again the BICYCLE example. Whether or not the hypothesis that *my bicycle has just been stolen* is effective at challenging my putative knowledge that my bicycle is where I left it depends on the plausibility of denying that I know that it is

false. In some situations, it will be plausible that I do know that my bicycle has not just been stolen—if I have just seen it, for example. In other situations, however, it might be plausible that I do *not* know that this hypothesis is false. This hypothesis is *effective* at challenging my putative knowledge in the latter, but not the former case. This point will be relevant and discussed again later on in this chapter.

DeRose (1995, 18) has argued that sensitivity considerations can best explain why certain sceptical hypotheses are plausibly not known to be false. DeRose expresses reservations about what can be called the ‘standard’ view. The standard view holds that it is plausible to deny that anyone knows that sceptical hypotheses are false, because the scenarios they describe are in principle *subjectively indistinguishable* by anyone from any other scenario that might obtain. The problem with the standard view according to DeRose, is that:

there are plenty of other phrases that can be used plausibly to describe our apparently limited epistemic position with regard to effective skeptical hypotheses. All of the following descriptions... have some initial plausibility: I cannot rule it out, I do not know that it does not obtain (and do not know whether it obtains), I can’t discern that it does not obtain... and I can’t distinguish its obtaining from its not obtaining, and so on, and so forth. But citing one of these to explain the plausibility of another does not occasion even the slightest advance in our understanding. (DeRose, 1995, 16)

DeRose’s point here is, I think relatively simple. Whether one thinks that S’s failure to know that $\neg sh$ is explained by S’s inability to distinguish between actuality and sh , or vice versa, would suggest that one does not really seem to have explained much at all. In order to break the apparent explanatory “circle of all-too-closely related terms of epistemic appraisal” when accounting for the intuitive denial that S knows that a given sceptical hypothesis is false, DeRose (1995) appeals to the notion of *sensitivity*. DeRose observes that a familiar feature of many sceptical hypotheses, such as the hypothesis that one is a brain in a vat, is that if the scenarios described by the hypotheses obtained, then S would still continue to believe that the hypothesis is false. In this way, for example, my beliefs about whether I am a brain in vat are not plausibly thought to match up to the facts; I would continue to believe that I am *not* a brain in a vat even if I was one. Likewise, my belief that I am not the victim of a malignant demon who is subjecting me to a complete deception is similarly *not sensitive*; I would continue to believe this even if it was false. Beliefs that would continue to be held even in the closest possible

world in which they are false are *insensitive*, and do not intuitively count as knowledge according to DeRose.¹ The idea that sensitivity is a condition on knowledge could then be expressed in the following way:

Sensitivity. For all S, p , if S knows that p , then S does not believe that p in the nearest possible world or worlds where $\neg p$.²

DeRose's sensitivity based account, points to the truth of certain subjunctive conditionals to explain the intuitive plausibility of denials that sceptical hypotheses are known to be false. According to DeRose, it is intuitively plausible that S does not know that $\neg sh$ *because* for virtually all effective sceptical hypotheses sh , it is true that if it were the case that sh , then S would believe that $\neg sh$. DeRose clarifies this when he states:

[This] explanation, in terms of subjunctive conditionals, can explain the plausibility of the other ways we feel inclined to describe our seemingly limited epistemic situation vis-a-vis effective skeptical hypotheses... *Because* we would still believe they weren't [true] even if they were. (DeRose, 1995, 19)

DeRose then loosely proposes that a necessary condition on effective sceptical hypotheses is that they be the sorts of proposition that we do not sensitively believe to be false. In the next section, before critically evaluating the sensitivity account of sceptical hypotheses, I outline the first two of three conditions on sceptical hypotheses argued to be necessary by Beebe: the experiential constraints. In section 2.6, I then present a case which I argue represents a counter-example to the necessity of the sensitivity condition, and the experiential constraints on sceptical hypotheses. On the basis of the same case, in section 2.6, I outline and argue against the necessity of a third constraint proposed by Beebe: the indication constraint.³

¹ (DeRose, 1995, 18).

² Adapted from (Pritchard, 2005, 48). The history of sensitivity considerations goes back to (Nozick, 1981).

³ The discussion focuses on the issue of the *effectiveness* of sceptical hypotheses—what makes them apparently capable of challenging our empirical knowledge. The issue of what distinguishes sceptical hypotheses from their counter-parts such as the stolen bicycle hypothesis is discussed subsequently to this.

2.3 Experiential constraint.

In this section, I outline Beebe’s experiential constraints. The first of these is the following, which Beebe takes to represent a necessary condition on being a sceptical hypothesis:

Experiential Constraint 1 (EC1): For a sceptical hypotheses *sh* to raise a significant sceptical challenge to *S*’s putative knowledge that *p*, it must be experientially possible for *sh* to be true. (Beebe, 2010, 466)

Beebe explains the notion of ‘experiential possibility’ in terms of Lewis’s notion of centred worlds in the following way:

w is a[n experiential] possibility for *S* iff in *w* *S*’s experiences and memories match *S*’s experiences and memories in actuality[.] (Beebe, 2010, 466)

Beebe claims that “experientially possible worlds for *S* will be subjectively indistinguishable from the world *S* actually inhabits” (ibid). Beebe takes it to underpin the “heart of the challenge... [which] is that: the sceptic alleges that my evidence is insufficient to tell me whether I am in a *Y*-world or a *Z*-world” (Beebe, 2010, 468). This constraint appears to be well suited to the sceptic’s purposes of challenging our claims to know anything about the external world. In particular, it highlights that sceptical challenges usually involve calling into question whether our sensory experiences allow us to know external world propositions. For example, something akin to the subjective indistinguishability of sceptical possibilities from everyday possibilities is central to Barry Stroud’s assessment of sceptical challenges:

There have been many versions of that fundamental idea. But whether it is expressed in terms of “ideas” or “experiences” or “sense data” or “appearances”... or whatever it might be, the basic idea could be put by saying... whatever it is that we get through that source of knowledge... whatever they might be that serve as the sensory “basis” of our knowledge, it does not follow that something we believe about the world is true. (Stroud, 1984, 549)

This assessment is echoed by Michael Williams, who concedes that this is:

a perfectly representative account of the main obstacle to knowledge of the world: ...it seems very hard to deny. At least, if typical sceptical counter-possibilities —I am dreaming right now (or have always been dreaming), I

am a brain in a vat, etc. — are even coherent, they seem to establish... the simple logical point that our experience could be just what it is and *all our beliefs about the world could be false*. (Williams, 1991, 74)

In addition Beebe proposes the following similar constraint on sceptical hypotheses:

Experiential Constraint 2 (EC2): For a sceptical hypotheses *SK* to raise a significant sceptical challenge to *S*'s putative knowledge that *p*, it must be experientially possible for *p* to be false. (Beebe, 2010, 466)

(EC2) differs from (EC1) in that the former says that the falsity of various external world propositions must be experientially possible for us, whilst the latter holds that the hypothesis itself must be experientially possible for us, in order for it to raise a significant sceptical challenge. According to (EC1), it is a necessary condition on a sceptical hypothesis *sh* that, if *sh* were true, our memories and experiences would be the same as they actually are. According to (EC2), it is also a necessary condition on a sceptical hypothesis *sh* that our memories and experiences can be the same as they actually are, and external world propositions be false. For my purposes here it will not be important whether there is more to be said about how (EC1) and (EC2) are similar, or differ. It will suffice to note that Beebe maintains each to be a necessary condition on sceptical hypotheses, and that this view is argued against below.⁴

With respect to (EC1), it is generally recognised that if it were the case that either the dreaming hypothesis or the *brain in a vat* hypothesis were true, then *S*'s experiences and memories would match *S*'s experiences and memories in actuality. For example, speaking of sceptical hypotheses such as the 'Evil Deceiver' and a version of the *brain in a vat* hypothesis, Michael Williams stresses that these sceptical hypotheses "are designed to accommodate any experiential data that we might cite in support of our ordinary" knowledge about the external world.⁵ Whilst Ernest Sosa makes a point of recognising that if "while dreaming we have real beliefs based on real phenomenal experiences, then a normal perceptual judgement could always be matched by a subjectively similar, similarly based judgement, made while one dreams" (Sosa, 2009, 2). Both sorts of hypotheses satisfy this constraint on being an effective sceptical hypothesis.

Let's call the two constraints (EC1) and (EC2) the *experiential* constraints. I take it that the experiential constraints are supposed to help provide some account of what

⁴ See (Beebe, 2010, 466, 470) for his position.

⁵ (Williams, 1991, 79).

2.4. THE CASE AGAINST THE EXPERIENTIAL, AND SENSITIVITY CONSTRAINTS.

constitutes a sceptical hypothesis. To investigate whether or not they do, it will be instructive to ask the following question. What does *sh*'s being experientially possible for *S* explain about why *sh* is capable of raising a challenge to knowledge of a empirical propositions? The idea, here, might be that when *sh* is experientially possible for *S*, there will be nothing in how the world appears to *S* that they can bring to bear on the question of whether *sh* is true or false. Pritchard explains the connection in following way:

The problem that a sceptical hypothesis... poses is that it is subjectively indistinguishable to the agent... And given this subjective indistinguishability, it seems that there will be nothing cognitively available to the agent which will suffice to indicate to her that she is in the non-sceptical scenario as opposed to the sceptical scenario. So... sceptical error-possibilities pose an *in principle* difficulty for the epistemic status of one's beliefs. (Pritchard, 2005, 24)

There is an apparent explanatory link here between the experiential constraints and the plausibility of denying that *S* knows that $\neg sh$. The idea here is that, we can explain why it is plausible to deny that *S* knows that $\neg sh$ by pointing out that it is experientially possible for *S* that *sh* (and subsequently— that $\neg p$). The problem with this explanation is that a hypothesis that it *not* experientially possible for *S* may still generate an intuitively plausible denial that *S* knows the hypothesis to be false; at least, this is what I argue for in the following section.

2.4 The case against the experiential, and sensitivity constraints.

I argue that neither the experiential constraints, nor the sensitivity condition are necessary conditions on a hypothesis raising a significant challenge to our external world knowledge. I think that the following case provides an example of a hypothesis that is used to raise a significant challenge to a subject's putative knowledge, but which neither meets the experiential constraints, nor satisfies the sensitivity condition. I argue that, insofar as this hypothesis is an analogue of any sceptical hypothesis, it represents a counter-example to the claim that in order to raise a sceptical challenge, it is necessary for a hypothesis to meet either the experiential constraints, or the sensitivity condition.

Before presenting the case, it will be helpful to comment briefly on what I do *not* take the case to show. In particular, I do not think that the case motivates any sceptical

conclusions about knowledge. The case is constructed so as to describe a situation in which various moves are made which correspond to those involved in standard sceptical arguments. However, the apparent cogency of the sceptical moves within the case do not demonstrate a sceptical conclusion *about* the case; no more so than that the apparent cogency of sceptical arguments demonstrates the truth of scepticism. The case is supposed to be a model of sceptical challenges. The purpose of the model is to make it easier to diagnose what is, and is not, going on within these challenges. For my purposes here, I maintain simply that what *is going on* is an analogue of a sceptical challenge involving some hypothesis. I contend, however, that neither the the experiential constraints, or sensitivity condition are met in this example.

Ellie is a new visitor at Jurassic Park, where she is in the Park's emergency bunker with the park's proprietor, Hammond, and his associate, Malcolm. Ellie is trying to find out where her colleague Alan is within the park. Within the bunker there is a single monitor for a CCTV surveillance system that shows the entire paddock. Due to frequent and violent storms at the park, there are constant faults in the surveillance system. Nonetheless, the surveillance system (including the monitor) has a special new feature called 'Doxasafe': the monitor shows a solid bright red screen if and only if there is any fault—for whatever reason—anywhere in the system. In fact, it was a legal condition of the Park that Doxasafe monitors were installed: if the system had not had Doxasafe, there would have been no Jurassic Park. However, Ellie is completely unaware of either the fact that there are regular faults, or that the system has Doxasafe. At this point, the system is not faulty and the monitor shows a live picture of the paddock, with Alan nowhere in sight. On the basis of looking at the monitor Ellie forms both the beliefs that the surveillance system is not faulty, and that Alan is not in the paddock. In fact, both of Ellie's beliefs are true. Whilst observing Ellie do this, Hammond says to Malcolm "She doesn't know that the monitor isn't faulty." Malcolm then responds, after reflecting for a moment, by saying "So, she doesn't know that Alan's not in the paddock either!"

Notice that both of Ellie's beliefs in this case are *sensitive* in the sense that the following relevant subjunctive conditionals are *false*. If the system were faulty, then Ellie would believe it was not, and, if Alan was in the paddock, then Ellie would believe that he was not.

2.4. THE CASE AGAINST THE EXPERIENTIAL, AND SENSITIVITY CONSTRAINTS.

In order for Ellie's belief that the monitor is not faulty to be sensitive, she must *not* believe that the monitor is not faulty in the closest possible world in which it *is* faulty. And in the closest possible world in which the monitor is faulty, it is clear that Ellie will *not* believe that the monitor is not faulty. The closest world in which the monitor is faulty is one in which the monitor shows a solid bright red screen, and it is plausible to think that Ellie will *not*, by looking at it, believe that the monitor is not faulty. Her belief that the monitor is not faulty, then, is sensitive. Likewise, Ellie's belief that Alan is not in the paddock is sensitive only if she does *not* believe that Alan is not in the paddock in the closest possible world in which he *is* there. In the closest possible world in which Alan *is* in the paddock, it is clear that Ellie will *not* believe that he is not in the paddock. The closest world in which Alan *is* in the paddock is one in which the picture on the monitor screen will show him there. So, I think it is plausible to say that *were* Alan in the paddock, then Ellie *would not* form a belief that he is not there on the basis of looking at this screen. Her belief that Alan is not in the paddock, then, is sensitive.

Imagine that Hammond is aware of both the facts that the monitors have this special feature and that Ellie is unaware of this feature, neither of which is implausible. Now, it is my contention that Hammond has said something that has a great deal of intuitive plausibility when he says to Malcolm that Ellie does not *know* that the monitor is not faulty. What Hammond would be doing by uttering that sentence is—at least on the face of it—denying that Ellie knows some proposition to be false, or situation not to obtain, etc.⁶ In so much as Ellie has a sensitive belief that the monitor is not faulty, the apparent fact that the faulty-monitor hypothesis can be used to raise a significant challenge to her putative knowledge that Alan is not in the paddock shows that it is not a necessary condition on doing so that Ellie insensitively believe that the hypothesis is *false*.

This case provides reason to think that a hypothesis can be capable of raising a significant sceptical challenge even when a belief that this hypothesis *is false* is a sensitive belief. Before outlining how the case provides a counter-example to Beebe's experiential constraints, it will be worth anticipating some potential objections to my argument here. It might be objected that the hypothesis that *the monitor is faulty* is not really a sceptical hypothesis. Perhaps for the reason that it seems too innocuous to pose a threat to Ellie's knowledge; for example, it does not entail the falsity of what she believes about Alan's whereabouts. I think that this is unconvincing. To see why, imagine again that Hammond tells Malcolm that Ellie does not know that the monitor is not faulty,

⁶ There is no obvious reason not to think this is exactly what Hammond is doing and intends to do by uttering this sentence. I think we should take Hammond's intentions at face value here.

and further points out to Malcolm that if Ellie knows that Alan is not in the paddock then she knows that the monitor is not faulty. Malcolm can then infer that Ellie does not *know* that Alan is not in the paddock. After all, the only basis on which Ellie believes Alan is not there is having looked at the monitor; so if Malcolm is not prepared to judge that Ellie even knows whether the monitor is accurately showing the paddock, then he is in a position to judge that she does not know he is not there.

It might also be objected that the faulty-monitor hypothesis is not really an effective sceptical hypothesis, for the reason that it is not capable of raising a challenge to Ellie's (or anyone else's) knowledge of a range of empirical propositions. In this way, the faulty-monitor hypothesis could be regarded as akin to the non-sceptical hypothesis that *my bicycle has just been stolen*. As such, it could be objected that the case does not provide a counter-example to the necessity of meeting the experiential constraints, or sensitivity condition on being an effective sceptical hypothesis. This is because the imagined hypothesis is not, 'strictly speaking', sceptical at all. I argue, however, that this objection fails; the faulty-monitor hypothesis is clearly effective in the sense identified above. It seems to pose a threat to an obvious *range* of propositions—namely, those that might be believed on the basis of looking at the monitor. The range of beliefs that the faulty-monitor hypothesis is capable of raising an epistemic challenge to is clearly *narrower* than a full-blown sceptical hypothesis. What I take away from this point, however, is that effective sceptical challenges come in degrees. Insofar as the faulty-monitor hypothesis calls into question a belief forming process, it is similar to the sceptical *brain in a vat*, and *dreaming* hypotheses, but not the non-sceptical *stolen bicycle* hypothesis.⁷ Moreover, the faulty-monitor hypothesis is arguably an *effective* sceptical hypothesis in the sense that it limits our epistemic position with respect to the challenged range of propositions. It is my contention that this is intuitively the case with respect to Hammond's imagined challenge to Ellie's epistemic position with respect to any proposition she believes on the evidential basis of looking at the monitor screen.

Finally, one last anticipatory point. One might worry that the case, if taken to be a genuine analogue of sceptical challenges, places too strong a requirement on perceptual knowledge, and subsequently would motivate sceptical conclusions. It is important to stress that I do not take the case to show that Ellie *does not* in fact know that faulty-monitor hypothesis is false, or that she does not *in fact* know that Alan is not in the

⁷ My view is that the stolen bicycle hypothesis has a maximally narrow scope, and does not obviously challenge a range of beliefs. Compare, for example, Dretske's (1970) cleverly disguised mule hypothesis. This differs from the stolen bicycle hypothesis in that the former is likely to raise broader sceptical doubts about one's evidential situation, in a way that the latter is not.

2.4. THE CASE AGAINST THE EXPERIENTIAL, AND SENSITIVITY CONSTRAINTS.

paddock. In just the same way that I do not take sceptical arguments to show that I do not know sceptical hypotheses are false, or know external world propositions. The fact that sceptical arguments are paradoxical is not, on my view, evidence for the incoherence of our concept of knowledge. Rather, it is a sort of evidence that sceptical arguments mislead us in ways that are not readily apparent.⁸ On my view, Hammond's *denial* that Ellie knows that monitor is not faulty, and Malcolm's denial that Ellie knows Alan is not in the paddock are both erroneous in the same sense, respectively, as the claims that 'I do not know that I am not dreaming', and that 'I do not know that I have hands' are also erroneous. And my contention is simply that the role played by the hypothesis is not fundamentally different between the case of paradoxical sceptical arguments, and the case of Ellie.

Consideration of the Ellie case provides a counterexample to Beebe's claim that (EC2) and (EC1) are necessary conditions on *effective* sceptical hypotheses. In this way, I think the case also allows us to provide a more in-depth explanation of DeRose's charge of explanatory circularity levelled against the account of sceptical hypotheses in terms of experiential possibility, or the subjective indistinguishability. Central to Hammond's imagined challenge to Ellie's putative knowledge that *Alan is not in the paddock* is the possibility that the monitor Ellie has looked at *is faulty*. It is this possibility that Hammond *denies* Ellie knows *not to obtain*, and so it can be thought to play the same role in Hammond's challenge as a sceptical hypothesis plays in a sceptical argument. It has been argued that the faulty-monitor hypothesis is indeed an effective hypothesis, since where S names Ellie, *sh*= that *the monitor is faulty* and *p*= that *Alan is not in the paddock*, the following two claims are intuitively plausible:

(i) S does not know that not-*sh*.

(ii) If S knows that *p* on the basis of looking at the monitor, then S knows that not-*sh*.

The hypothesis in this case, however, is not obviously experientially possible for S in either scenario. If it were the case that the monitor was faulty, then Ellie's experiences and memories would not match her experiences and memories in actuality. If the monitor was faulty, then, so long as she looks at the monitor, Ellie would have visual experiences of a solid bright red screen. This experience does not match anything that Ellie experiences

⁸ See, for example, chapter 6, where I propose an account of the erroneous plausibility of sceptical denials as a sort of 'cognitive illusion'. Compare 4 and 5, where I discuss and critique approaches to resolving sceptical paradoxes in terms of semantic context-sensitivity.

in the actuality of the case. This demonstrates that it is not experientially possible for Ellie that *the monitor is faulty* is true. So, the plausibility of a denial that Ellie knows that this faulty-monitor hypothesis is false cannot be explained by its being experientially possible for Ellie. It might be objected on behalf of the experiential constraints and their explanatory role vis-a-vis sceptical denials, that it is not plausible that in this case it is experientially possible for Ellie that *Alan is not in the paddock* is false. Subsequently, by (EC2), the hypothesis that *the monitor is faulty* is not a genuine sceptical hypothesis capable of raising a significant challenge to Ellie's putative knowledge of *that* fact.

It does seem as though Ellie's experiences and memories would *not* match those in actuality were it the case that Alan was in the paddock; some representation of him would appear on the screen she is looking at, for example. Therefore, it is not experientially possible for Ellie that her belief concerning Alan's whereabouts is false. But this is no objection. The idea, here, is that the two following claims about the case have a great deal of intuitive plausibility:

(i*) Ellie does not know that the monitor is not faulty on the basis of looking at it.

(ii*) If Ellie knows that Alan is not in the paddock on the basis of looking at the monitor, then Ellie knows that the monitor is not faulty on the basis of looking at it.

And this is all that one needs, on my view, in order to effectively challenge Ellie's putative knowledge that *Alan is not in the paddock*. The hypothesis that *the monitor is faulty* is, then, *effective* at raising a significant challenge, in this case, with respect to Ellie and the putatively known proposition that *Alan is not in the paddock*. However, it is neither experientially possible for Ellie that the hypothesis is true, nor that the proposition is false. So, the *faulty-monitor* hypothesis does not satisfy either (EC1), or (EC2). Yet I argued that the *faulty-monitor* hypothesis *is effective*, in the sense of raising a significant challenge to Ellie's putative knowledge of some external world proposition. So, I conclude, (EC2) and (EC1) are not necessary conditions on being an effective sceptical hypothesis.

The case presented here reveals that a significant challenge to S's putative knowledge of some external world proposition *p* can be raised by means of a hypothesis that: does *not* entail the falsity of *p*, *is* sensitively believed to be false by S, and *is* subjectively distinguishable from actuality by S. In the next section, I critically assess a further potential necessary constraint on effective sceptical hypotheses, also endorsed by Beebe.

2.5 Indication constraint.

In the previous section, I argued that the presented case provided a counter-example to the claim that the sensitivity condition, and the experiential constraints are necessary conditions on being a sceptical hypothesis. In this section, I outline and critically examine a third constraint put forward by Beebe as a necessary condition on sceptical hypotheses: the indication constraint. I argue that it too can be rejected on the basis of the case above. Consider the following constraint on sceptical hypotheses proposed by Beebe:

Indication Constraint (IC): In order for a sceptical hypothesis *SK* to raise a significant sceptical challenge to *S*'s putative knowledge that *p*, *SK* must indicate how *S* could believe that *p* on the basis of *S*'s evidence and *yet not know that p*.⁹

Beebe provides some brief considerations in support of the *indication* constraint. The argument begins by noting that only the third but not the first two of the following hypotheses can raise a significant *sceptical* challenge to *S*'s putative knowledge that *the animal in the pen is a zebra*:¹⁰

(B1) The animal in the pen is a lion.

(B2) The animal in the pen is not a zebra.

(B3) The animal in the pen is a mule cleverly disguised to look like a zebra.

This seems correct. Next Beebe points out that a significant difference between (B3) and (B1)-(B2) is that, when used to challenge some subject *S*'s putative knowledge that *the animal in the pen is a zebra*, only the former indicates how it is that *S* can believe what they do (i.e. that it's a zebra) on the basis of their evidence and yet not *know* this. Beebe concludes on this basis that (IC) is correct. This consideration appears to be compelling.

There are other aspects, however, which Beebe does not comment on not; ones that I think, ultimately, undermine his point. For example, there are a number of other finer-grained differences between (B1), (B2) and (B3) that are not explicitly mentioned in this argument. Before discussing these differences, however, an important similarity should also be highlighted. The similarity consists in the fact that it seems plausible to say that *S* not knowing any one of these supposed sceptical hypotheses to be false is incompatible with *S* knowing that *the animal in the pen is a zebra*. In other words, each

⁹ (Beebe, 2010, 453).

¹⁰ (Beebe, 2010, 452); cf. (Dretske, 1970).

satisfies one aspect of being an *effective* hypothesis in the sense identified here: viz. that it is intuitively plausible that someone knows the hypothesis to be false, *when they know* that the animal in the pen is a zebra. This indicates, I think, that the relevant aspect of ‘raising a significant sceptical challenge to S’s putative knowledge that *p*’ is here, on Beebe’s view, the intuitive plausibility of denying that the hypothesis is known to be false.

Let’s turn now to the differences. Firstly, it is possible to recognise that (B1) differs from both (B2) and (B3) in one important respect: *experiential possibility*. That is, (B1) will already count as an *ineffective* sceptical hypothesis on Beebe’s account in virtue of failing to satisfy (EC1). Beebe argues that (B3) *is*, but (B2) is *not* an effective sceptical hypothesis. The former *effective* hypothesis satisfies the indication constraint: it indicates how S could believe that *p* on the same basis as in actuality and yet not know that *p*.¹¹ So, Beebe concludes, the indication constraint is a necessary constraint on raising an effective sceptical challenge to S’s putative knowledge that *p*.

To draw out the point here, it is instructive to briefly compare the indication constraint with the sensitivity account of effective sceptical hypotheses discussed above. There are, it seems, plenty of hypotheses which we think both are known to be false, and should be if we know some external world proposition, but which we also recognise can only be insensitively believed to be false. Claims of the form of (i)—‘S does not know that not-*sh*’—are unlikely to be intuitively plausible, for example, when *sh* is the hypothesis that S is ‘an intelligent dog who is always incorrectly believing’ that they have hands.¹² The reason this hypothesis fails to be effective, according to DeRose, is that it does not explain or indicate why S holds the false belief that is attributed to them in the scenario it describes.¹³

In order to assess whether the indication constraint helps us to understand effective sceptical hypotheses, it will be instructive to check whether this constraint is satisfied in the case that I introduced in section 2.6. It is not immediately obvious that the hypothesis in this case (that *the monitor is faulty*) does satisfy the indication constraint. In order to do so, the *faulty-monitor* hypothesis would have to indicate to us how it is that Ellie could believe on the basis of her evidence that *Alan is not in the paddock*, and yet not know

¹¹ For Beebe, the basis of S’s belief that *p* is S’s evidence: I will simply refer more generally to S’s basis for belief.

¹² (DeRose, 1995, 22). One might also say, as DeRose does, that it seems to us—to the contrary—as though S *does in fact* know this sort of hypothesis to be false. I do not disagree with DeRose on this point, although for present purposes it is only important that the denial of that this is known is itself not plausible.

¹³ (DeRose, 1995, 23).

that he is not. However, I do not think that the hypothesis that *the monitor is faulty* does this. The hypothesis does not, for example, tell us that *were* the monitor faulty then Ellie *would* believe what she does about Alan's whereabouts, and yet not know the same thing. To see this, let us recall what I take to be the point of Hammond's challenge involving this hypothesis. This is that, in virtue of it not being a normal monitor, if the monitor had been faulty then Ellie would not believe what she does about Alan's whereabouts at all. In this way, it does not seem as though the hypothesis indicates anything over and above a literal falsehood about the monitor: that it is faulty.

Perhaps this is not the best way to understand how the indication constraint might apply to this hypothesis. Crucial to Hammond's challenge is the collateral information that Ellie's beliefs about what the monitor would be like *if* it had been faulty are themselves erroneous. Without this information it is not clear how one could go about judging whether she knows. It is within the context of *this* information that the hypothesis seems, *at most*, to indirectly convey to us how it is that Ellie could believe, *but not know* that *Alan is not in the paddock*, on the basis of her evidence. It seems to reveal little more to us than that Ellie *actually does* believe that Alan is not in the paddock on the basis of having looked at the monitor, and yet may not *actually* know that he is not there.

I think there are further explanatory issues with the indication constraint, irrespective of whether the hypothesis that *the monitor is faulty* plausibly satisfies this constraint directly, or indirectly. It is not clear that the plausibility of Hammond's denial that Ellie knows that the monitor is not faulty relies in any way on considerations of Ellie's belief about Alan's whereabouts. The plausibility of this denial would not be diminished if (*ceteris paribus*) Alan was in the bunker too and nobody there was wondering where he was! Ellie is plausibly ignorant of the fact that the monitor is not faulty independently of any other beliefs that she arrives at by means of the monitor. This suggests, *pace* Beebe and DeRose, that whether or not it is plausible that S does not know that $\neg sh$ is independent of whether *sh* indicates (either directly or indirectly) how S could believe that *p* on the same basis as in actuality, and yet not know that *p*.

There may be more to an effective sceptical hypothesis *sh* being capable of raising a significant challenge to S's putative knowledge that *p* than simply S plausibly not knowing that $\neg sh$. It must also be plausible that if S knows that *p* then S knows that $\neg sh$. It is in this aspect of being an effective sceptical hypothesis then that we find a definite link between considerations of S's belief that *p* and *sh*, which is something that the indication constraint attempts to capture. So, perhaps an examination of what the indication constraint contributes to explaining the plausibility of claims of the form of

the second sceptical premise (i.e. ‘if S knows that p , then S knows that not- sh ’) will reveal more. I explore this next.

2.6 Bad possibilities.

The foregoing argument has been that neither the experiential, nor the indication constraints put forward by Beebe should be considered necessary conditions on being a sceptical hypothesis. I now aim to put forward a positive proposal about what could be said to be necessary for being an effective sceptical hypothesis. In developing this proposal, I will draw again on considerations already highlighted in chapter 1, concerning the notion of ‘bad’ possibilities, or scenarios. In this section, I set out the parallels between the indication constraint and characterisations of effective sceptical hypotheses as ‘bad’ possibilities. After pointing towards the limitations of understanding effective sceptical hypotheses in these terms, I use them, in section 2.7, as a springboard for proposing an alternative account.

I think that consideration of the indication constraint could help to explain why claims of the form ‘if S knows that p , then S knows that $\neg sh$ ’ are plausible for effective sceptical hypotheses. The indication constraint says that it is a necessary condition on any sceptical hypothesis sh being effective at challenging S ’s putative knowledge that p , that sh be a scenario in which S believes, but does not know that p on the basis of their evidence. It is not surprising that virtually all paradigmatic effective sceptical hypotheses fit this pattern: a brain in a vat does not know it has hands, an individual looking at a cleverly disguised mule does not know it is a zebra, and the victim of a complete perceptual deception knows nothing about the world—at least, not on the basis of perception. This might be thought to be obvious given the closure principle, the factivity of knowledge, and the fact that, in each of the aforementioned examples of sceptical hypotheses, these hypotheses entail the falsity of external world propositions. This, however, is not the case with respect to many other well-recognised hypotheses capable of raising a significant challenge to our putative knowledge. The *dreaming* hypothesis, for example, does not entail the falsity of many external world propositions, but is capable of raising a sceptical challenge to my putative knowledge of them. Likewise, the hypothesis that *half of the animals in the pen are cleverly disguised mules*, does not entail that *this animal in the pen is a zebra* is false, but it is capable of raising a sceptical challenge to my putative knowledge that this animal in the pen is a zebra.¹⁴

¹⁴ See (McGrath, 2013) for discussion of this hypothesis.

I take it then that an explanatory advantage of the indication constraint, even though it does not quite deal well with my example of the *faulty-monitor* effective hypothesis, is that it goes some way towards providing a more general account of sceptical hypotheses. It applies in equal measure to hypotheses that are not compatible with the targeted proposition p , such as the *brain in a vat* hypothesis, as it does to those that are compatible with p , such as the *dreaming hypothesis*. Insofar as the indication constraint is to be rejected on the grounds that it fails to apply to the case of Ellie, this unifying aspect of it will nonetheless be a desirable feature in any alternative account.

In seeking an account of effective sceptical hypotheses, I have suggested that more generality is preferable to less. Specifically, I take it that it is a desirable feature of an account that it both apply equally to compatible and incompatible hypotheses, in the sense outlined above, and apply to my Ellie case. In seeking an account like this, I think it will be instructive to reconsider some influential characterisations of effective sceptical hypotheses that fall short of having the desired generality. I think that there may be some valuable lessons to be gained from these characterisations, despite their limitations, that can help to shed more light on what an effective sceptical hypothesis must do.

The first account I am interested in considering has been offered by Williamson (2000), who characterises effective sceptical hypotheses as ‘bad’ possibilities in the following way:

In the good case, *things appear* generally as they ordinarily do, *and are that way*; one believes one proposition p (for example, that one has hands), and p is true; by ordinary standards, one knows p . In the bad case, things *still appear* generally as they ordinarily do, *but are some other way*; one still believes p , but p is false; by any standards, one fails to know p , for only true propositions are known. (Williamson, 2000, 165, my emphasis.)

In the bad scenarios, described by effective sceptical hypotheses, S still believes that p on the basis of it seeming to them that p , but S does not know that p because p is false. But what is it about these scenarios that is supposed to be epistemically bad? If there can be knowledge of contingent propositions at all, then surely there will be possibilities in which these propositions are false, and thereby not known. Moreover, it has been continuously stressed here that in order to be effective—at raising a significant challenge to S ’s putative knowledge of some external world proposition p —it is not necessary for a hypothesis to entail the falsity of p . It appears that the list of effective sceptical

hypotheses with respect to p is not exhausted by the list of bad possibilities with respect to p as recognised by Williamson. Specifically, effective sceptical hypotheses such as the *dreaming* hypothesis, and the weaker *half-fakes* hypothesis, are not obviously examples of Williamson's bad possibilities. I submit then that an account of effective sceptical hypotheses in terms of Williamson's notion of bad possibilities does not provide us with the general account we are seeking.

In contrast to Williamson, Pryor (2000) recognises that not all effective sceptical hypotheses entail the falsity of putatively known external world propositions. And it is from Pryor's recognition of this fact, and other considerations concerning effective sceptical hypotheses that I now take my cue in pursuing an alternative account. In particular, I take Pryor's useful insight, with respect to sceptical hypotheses, to be that certain possibilities are not bad *simpliciter* but only relative to the propositional content and grounds for—or experiential basis of—a perceptual belief that p .¹⁵ In line with the recognition that effective sceptical hypotheses may either entail the falsity of the putatively known proposition or not, Pryor outlines two distinct ways in which a sceptical hypothesis sh may describe a bad possibility q with respect to a perceptual belief that p :

Say that some grounds E you have [for a belief that p] “allow” a possibility q iff the following counterfactual is true: if q obtained, you would still possess the same grounds E ... So, we might want to count a hypothesis as “*bad*” for the purposes of a skeptical argument just in case it is (and is recognized to be) incompatible with what you purport to know [p], but it is nonetheless “*allowed*” by your grounds E ..., [or], if it could undermine your experiences. (Pryor, 2000, 527)

On this view, a hypothesis sh will be effective at raising a significant sceptical challenge to S 's putative knowledge that p only when it describes a bad possibility with respect to p and S 's grounds, E , for believing that p . A ‘bad’ possibility, in this sense, is one in which S believes that p on the same grounds E as they do in actuality, but where either p is false, or else E no longer provides any justificatory support for S 's belief that p .¹⁶ Since S does not know that p if either p is false, or S 's belief that p lacks justificatory

¹⁵ (Pryor, 2000, 528).

¹⁶ See (Willenken, 2011, pp.6-7) for an explanation of the function of an ‘underminer’—Pryor's term for a hypothesis that could undermine your experiences. Pryor maintains certain perceptual beliefs about the world are based on grounds or experience, which are what justify these beliefs. Others maintain that beliefs are based upon and justified by evidence, either propositional or non-propositional. For present purposes I remain neutral on whether beliefs are best described as being based on or justified by evidence, experience, grounds, or reasons; requiring only that beliefs are based in some way.

grounds, it is clear that a common feature of bad possibilities with respect to p and E is that, *in these possibilities*, S does not know that p on the basis of E .

In this way, Beebe's indication constraint has clear parallels with this characterisation of sceptical hypotheses in terms of bad possibilities. Bad possibilities are those in which a subject could believe, but not know, a proposition p on some evidential basis, on which they would ordinarily base a belief that p . The close connection between the indication constraint and Pryor's characterisation of sceptical hypotheses as 'bad' possibilities should not be surprising. After all, Pryor explicitly states that he considers a possibility to be bad "just in case it has the special features that characterize the skeptic's scenarios—whatever those features turn out to be" (Pryor, 2000, 527). If being a sceptical hypothesis simply amounted to describing a 'bad' possibility, where a 'bad' possibility is just one that has those features that allow it to serve the requisite role in a potentially threatening sceptical argument, then it remains unclear whether the notion of 'badness' can help us meaningfully distinguish between sceptical hypotheses that are *effective* from those that are not.

Presumably, then, the 'badness' of a possibility consists in it satisfying the necessary conditions on being an effective sceptical hypothesis. I think this is a plausible way to interpret Pryor's view of bad possibilities and effective sceptical hypotheses. So, perhaps the indication constraint is a necessary condition in the following sense: a hypothesis sh is effective with respect to S 's putative knowledge that p when it describes a bad possibility, where this notion is explicated in the following way:

a possibility is bad, with respect to an experience and a proposition, iff it is a possibility in which one would have the experience and ordinarily rely on it in believing the proposition but in which one would not thereby gain perceptual knowledge of the proposition. (McGrath, 2013, 535)

This latter-most characterisation of a 'bad' possibility, together with the closely related indication constraint, seem to provide some insights into what features of effective sceptical hypotheses are doing the necessary work in challenging our putative knowledge of empirical propositions. I argue, however, that neither provide a satisfactory account of the necessary conditions on effective sceptical hypotheses; neither, I have argued, applies to the example of the effective hypothesis involved in the Ellie case. In what follows, I suggest a weaker alternative constraint on effective sceptical hypotheses, which I argue succeeds where the others have failed in applying to the Ellie case.

Consider again the novel Ellie case I introduced in section 2.4. I have argued that the case involves an effective epistemic challenge, which is an analogue of sceptical challenges. In this case, however, the hypothesis utilised in the raising of a significant challenge to the protagonist's putative knowledge of some proposition *does not* describe a possibility which fits with the above definition of a bad possibility. In any possibility in which the monitor is faulty, Ellie neither has the monitor-related experiences she actually does have, nor does she believe the relevant proposition. But it is obvious that there is some aspect of this hypothesis that is 'bad' in the sense that Hammond is able to utilise it effectively in raising an epistemic challenge to Ellie's belief that Alan is not in the paddock. This case is, however, importantly dissimilar to the various examples of sceptical challenges contained within the cases from recent epistemology which involve extraordinary possibilities (brains in vats, dreaming, or malignant demons). The difference is, as I have noted in chapter 1 and section 2.2, one of scope. I think that this difference in scope is an aspect of effective epistemic challenges that requires explanation. But I think we will be in the best position to give this explanation once we have a general account of the constraints on being an effective sceptical hypothesis, which applies no matter how narrow or wide the range of its epistemic challenge.

Within the next section, I develop an alternative explanation of 'bad' possibilities, in terms of the basis for a belief, and then go on to propose an account of the necessary conditions on an effective sceptical hypothesis capable of raising a significant challenge to knowledge of external world propositions. I argue that this constraint captures the common core shared by effective hypotheses, and can accommodate the differences between them in terms of scope.

I noted above that it is an advantage of the indication constraint that it accounted for both those effective sceptical hypotheses that entail the falsity of everyday empirical propositions, and those that do not. I think that this generality is desirable in an account, and should be sought. The indication constraint, however, is not satisfied by the faulty-monitor hypothesis. Instead, I argue that this hypothesis satisfies another necessary condition on effective sceptical hypotheses, and one that also generalises in the desired way. On my view, what makes the faulty-monitor hypothesis an effective sceptical hypothesis is that were it true, then it would be consistent with Ellie believing that Alan was not in the paddock *on the same evidential basis as she actually does* that this belief is false. After setting out my constraint, and applying it to various examples, I go on to argue that it is preferable to the rejected necessary conditions on effective sceptical hypotheses.

2.7 Proposed condition on sceptical hypotheses.

My aim now is to draw upon the above considerations to sketch a positive account of the necessary conditions on being an effective sceptical hypothesis—one capable of raising a significant challenge to our putative knowledge of a range of propositions. To draw this out, I think it will be instructive to explore the pertinent question here as to *why* it is that Ellie *would* not know that *if* that hypothesis were true. So, what explains Ellie's lack of knowledge in *that* scenario?

The explanation is not that the putatively known proposition would be false, since it could be true that Alan is not in the paddock even if the monitor is faulty. Neither would it be because Ellie's belief may now too easily have been false, even though the believed proposition may be true. This is a popular explanation of S's inability to know that *p* when the possibility described by *sh* obtains. Yet, this remains an entirely implausible explanation for the case being considered. Ellie does not form the belief that Alan is not in the paddock in any nearby world in which the monitor is faulty, and subsequently, does not have an unsafe belief that this is the case. What, then, is going on with the *faulty-monitor* hypothesis in the Ellie case? What is it that makes it epistemically bad in regards to her believing that Alan is not in the paddock? The basis for Ellie's actual belief about Alan's whereabouts is having looked at the monitor.¹⁷ It is undeniable that were the faulty-monitor hypothesis true, then it would be epistemically bad for Ellie to believe that Alan is not in the paddock on *that* basis. Looking at a faulty monitor is no way to know where someone might be. Ellie stands to acquire as much knowledge about Alan's whereabouts by looking at a faulty surveillance monitor as she does by simply guessing. Ellie does not know that Alan is not in the paddock, if she believes this proposition on the basis of either these two methods.

Let us say that a basis B for belief that *p* is 'bad' when forming a belief that *p* on the basis of B precludes knowing that *p*. But a basis B that is bad for a belief in one proposition will not by itself be bad for belief in another. Ellie's counting of visible tables, for example, would be a bad basis on which to believe something about Alan's whereabouts, but not a bad basis on which to form a belief about how many tables she can see. The badness of a basis B for belief is then relative to the propositional content of the belief. Since looking at the monitor is not obviously a bad basis for Hammond, for example, to believe that Alan is not in the paddock, the badness of a basis B for belief that *p* is then also relative to the background knowledge of the agent forming the belief

¹⁷ I take it this is uncontroversial. Put another way: if Ellie had not looked at the monitor, it seems as though Ellie would have no reason to believe what she does.

that p . So, S's basis B for a belief that p is bad when S's believing that p on the basis of B precludes S knowing that p .

On the basis of the above considerations, I propose then the following putative constraint on effective sceptical hypotheses:

Belief basis Constraint: A hypothesis sh is capable of raising a significant epistemic challenge to S's putative knowledge that p , *only if* were sh true, then not- p would be consistent with S believing that p on the same evidential basis as they actually do.

Before running through the advantages of my proposed constraint by applying it to the various examples of sceptical hypotheses above, it is worth comparing my proposal from Pryor's account of sceptical hypotheses. My proposed constraint is similar to the characterisation of 'bad' possibilities Pryor gives in the sense that it centers around the notion of a subject's evidential basis for believing a proposition. The difference consists in that my account, but not Pryor's, remains neutral on whether S forms a belief that p at all in the scenarios described by sceptical hypotheses. On Pryor's view, I take it, the idea is that if an effective sceptical hypothesis sh is true, then a subject S will believe that p on the same grounds E that they would ordinarily do if sh were false (e.g. their experiences). Furthermore, if an effective sceptical hypothesis is true, then *in that case* S would not know that p because either, depending on the hypothesis, p is false, or E does not provide any epistemic support for p . My constraint, in contrast, is not committed to the idea that if an effective sceptical hypothesis is true, then a subject S will believe that p on the *same* grounds as they *would* do if the hypothesis were false. Instead, it is committed only to the weaker idea that if an effective sceptical hypothesis sh is true, then it is *possible* for S to falsely believe that p on the *same* evidential basis that they would otherwise do if sh were false.

In the remainder of this section, I will discuss how this constraint applies both to the various examples of recognised effective sceptical hypotheses considered above (i.e. the *brain in a vat*, *dreaming*, *cleverly disguised mules*, and *half-fakes* hypotheses), as well as to the novel example of the effective sceptical hypothesis presented in the Ellie case. Lastly, I suggest some ways in which the proposed necessary condition on sceptical hypotheses helps to explain the variable effectiveness of the non-sceptical *stolen bicycle* hypothesis.

With respect to the Ellie case and the *faulty-monitor* hypothesis, I argued that this was an example of an effective sceptical hypothesis. On my view, the *faulty-monitor*

hypothesis does satisfy the belief basis constraint. It satisfies this constraint only if it is the case that were the *faulty-monitor* hypothesis true, then it would be consistent with Ellie believing that Alan is not in the paddock on the *same* basis as she actually does that Alan *is* in the paddock. And this, I submit, is intuitively the case. Ellie's actual basis for believing that Alan is not in the paddock is looking at the monitor. Were the *faulty-monitor* hypothesis true, then the screen would show a solid red screen. In that case, if Ellie formed a belief that Alan is not in the paddock on the *same* basis as in the actual case, her belief would be formed on the basis of looking at the monitor which has a red screen. In which case, it would be consistent with Ellie believing that *Alan is not in the paddock on that basis*, that her belief is false (i.e. Alan *is* in the paddock). In this way, the faulty-monitor hypothesis satisfies my proposed constraint on effective sceptical hypotheses.

The *brain in a vat* and *cleverly disguised mules* hypotheses are obvious examples of effective sceptical hypotheses, albeit different in terms of their scope, and both satisfy the belief basis constraint. Each satisfies this constraint only if it is the case that were the hypotheses true, then it would be consistent with me, for example, believing the target propositions, respectively, that *I have hands*, and that *the animal in the pen is a zebra*, on the *same* basis as I actually do. And this, on my view, is plausibly the case. My actual evidential basis for believing that *I have hands* is plausibly that I have perceptual experiences as of hands. Were the *brain in a vat* hypothesis true, then plausibly my perceptual experiences would be indistinguishable from how they actually are, but it would be false that I have hands. As such, in that case, it would be consistent with me believing that I have hands on the *same* evidential basis as I otherwise would, and this belief be false. In this way, the *brain in a vat* hypothesis satisfies the belief basis constraint on effective sceptical hypotheses. Likewise, my actual evidential basis in Dretske-style cases might plausibly be my looking at the animal in the pen. And were the *cleverly disguised mules* true, then the animal would plausibly appear indistinguishable to me from a zebra, but it would be false that the animal in the pen is a zebra. It would then, in that case, be consistent with me forming a belief that the animal in the pen is a zebra on the same evidential basis as I actually do that this belief is false. So, the *cleverly disguised mules* hypothesis also satisfies the belief basis constraint.

The *dreaming* and *half-mules* hypotheses are plausible examples of effective sceptical hypotheses that differ in their scope, and both differ from the *brain in a vat* and *cleverly disguised mules* hypotheses insofar as they are not compatible, respectively, with the propositions that *I have hands* and that *the animals in the pen is a zebra*. In each case,

the effective sceptical hypothesis too satisfies the belief basis constraint. In the former case, there *would* be no inconsistency between me believing the proposition that *I have hands on the same evidential basis as I actually do* and this proposition being false, *were it the case that the dreaming hypothesis is true*. Likewise, in the latter case, there *would* be no inconsistency between me believing the proposition that *the animal in the pen is a zebra on the same evidential basis as I actually do* and this proposition being false, *were it the case that the half-mules hypothesis is true*. So, again, both hypotheses satisfy my proposed constraint on effective sceptical hypotheses. One last point to consider is how my constraint applies to the examples of supposedly *ineffective* sceptical hypotheses considered in section 2.5: (B1) and (B2). These cases are interesting, and there are a couple of points worth making about them. Firstly, these hypotheses do seem to satisfy the belief basis constraint. If it *were* true either that *the animal in the pen is a lion* or that *the animal in the pen is not a zebra*, then it *would* be possible for me to *falsely* believe that it is a zebra on the basis of looking at the animal.¹⁸ My belief basis constraint is, in that sense, very weak. On my view, insofar as neither (B1) nor (B2) are plausibly thought of as being effective sceptical hypotheses, I think it is likely that they both fail to satisfy some further necessary condition.¹⁹

To contrast, let's consider the *stolen bicycle* hypothesis, which is not an example of a sceptical hypothesis. Whether or not it is effective at raising an epistemic challenge to my putative knowledge that *my bicycle is where I left it* can depend on the context. Insofar as this hypothesis is not sceptical, it is not a necessary constraint on it being effective that it satisfies the belief basis constraint. Nonetheless, I submit that consideration of this constraint can help to account for why it may be effective in some circumstances, but not others. The *stolen bicycle* hypothesis might plausibly be effective, for example, if the only evidential basis I actually have for my belief that my bicycle is where I left it is my memory of simply leaving it there 4 days ago. In this case, I submit that *were* the hypothesis *true*, it *would* be consistent with my believing that *my bicycle is where I left it* on the same evidential basis as I actually do that my belief is false. When my bicycle has been stolen, it is consistent with me believing on the evidential basis of remembering

¹⁸ This would be *possible*, strictly speaking, *even if* I had visual experiences as of looking at a lion, or a gazelle, etc. Of course, extra information would be required to *explain why* I would believe that an animal that was a zebra. But that does not mean that the indication constraint was correct. The argument that these fail to be effective sceptical hypotheses because they do not satisfy the indication constraint begs the question. The *faulty-monitor* hypothesis is plausibly effective, but likewise does not satisfy the indication constraint, as I showed in 2.6.

¹⁹ On my view, these hypotheses fail to be effective at raising a sceptical challenge to my putative knowledge that the animal in the pen is a zebra because they are not plausibly *epistemically possible* for me. I develop this line of thought further in chapter 3.

leaving it there 4 days ago that it *is not there*. Compare this to the example of the *cleverly disguised mule* hypothesis. This hypothesis is effective at raising a sceptical challenge because it calls into question my actual basis for believing that this animal in the pen is a zebra, namely, my background assumptions about the zoo, and it seeming to me perceptually as though the animal is a zebra.

Yet, the *stolen bicycle* hypothesis might plausibly *not* be effective, for example, if the evidential basis I actually have for my belief that my bicycle is where I left it is having *just* looked at it in the same place I left it. I propose that we might, now, plausibly explain this along the following lines. In *these* circumstances, it is not the case that *were* the *stolen bicycle* hypothesis true, it *would not* be consistent with my believing that my bicycle is where I left it on the *same* evidential basis as I *actually* do that this belief is false. In the closest world in which my bicycle has been stolen, it is difficult to understand *how it could* be the case both that I form a belief that it is where I left it on the basis of having just seen it there, and this belief be false. In order for the stolen bicycle hypothesis to be effective at raising an epistemic challenge something more needs to be going on. What would seem to be required is some further *sceptical* hypothesis that would support a sceptical attitude towards my actual basis for believing that my bicycle is where I left it.

That concludes the discussion of how the proposed belief basis constraint on effective sceptical hypotheses applies to the recognised examples of the such hypotheses, as well as the faulty-monitor example introduced in this chapter. In the next section, I put forward some brief comments by way of an anticipatory defence of the proposed account of effective sceptical hypotheses.

2.8 Concluding remarks.

In this chapter, I have been concerned with assessing accounts of effective sceptical hypotheses, namely, those capable of raising significant challenges to our putative knowledge of a range of external world propositions. I have assumed that a hypothesis *sh* is capable of ‘raising of a significant challenge to putative knowledge’ of some empirical proposition *p* if claims of the form (i) and (ii) are plausible. I argued that the Ellie case reveals that accounts of effective sceptical hypotheses in terms of the sensitivity condition, experiential constraints, and the indication constraint do not provide us with the desired understanding of what a hypothesis must to in order to effectively raise a sceptical challenge. While these accounts may offer some explanation in specific cases,

they do not present, as I have shown, any necessary condition on effective sceptical hypotheses. In particular, I argued that the case of Ellie reveals that these are not necessary conditions on sceptical hypotheses; an effective sceptical challenge can be raised *without* these conditions being met.

Instead, I drew on the following two considerations to motivate the novel belief basis constraint on effective sceptical hypotheses. First, in order to be effective at raising a significant challenge to S's putative knowledge that p , a sceptical hypothesis sh must describe a scenario *in which* S's actual basis B for their belief that p is bad. Second, a basis B for a belief that p formed by S is bad just in case $\neg p$ is consistent with S believing that p on the basis of B. On this basis, I argued that it is plausible that effective sceptical hypotheses must be such that, if they were true, then it would be consistent with a subject's believing the targeted propositions, on the same evidential basis as they actually do, that these propositions are false.

On my view, the belief basis constraint captures a necessary condition on sceptical hypotheses, and one that I think has considerable advantages over existing accounts. There is a strong case to be made for the constraint over those argued against, owing to its comparative success with respect to providing a more general account of effective sceptical hypotheses. This success was highlighted in the previous section, where I argued that the proposed constraint could be applied to give a unifying account of both p -compatible and p -incompatible examples of both maximally-wide scope, and limited-scope effective sceptical hypotheses, and the novel faulty-monitor example. Describing a scenario in which S has a bad basis for a belief that p is a necessary part of what it takes to raise a significant challenge to S's putative knowledge that p . I take it that my proposed account is at no disadvantage to the competitors considered here, given its comparative success in accounting for various recognised examples of effective sceptical hypotheses, and the novel example of the *faulty-monitor* hypothesis.

In the next chapter, I explore further the diagnostic work on sceptical arguments, and focus there on the epistemic principles underwriting their premises. This is a natural next step given the discussion of sceptical hypotheses given here. The following, for example, remains a pertinent question: why is it that Hammond's denial that Ellie knows that the monitor is not faulty is so plausible? The following are what I take to be some plausible explanations: she hasn't ruled out, or eliminated that the monitor is faulty; for all *she* knows, the monitor *is* faulty. Recall, however, the worry I highlighted in section 2.2, about the apparent explanatory vacuity of these responses. I discuss and develop these issues further, in the next chapter, to propose and defend a promising

account of the epistemic principle underwriting sceptical arguments from ignorance.

EPISTEMIC PRINCIPLES.

3.1 Introduction

In this chapter, as in the previous chapters, I am concerned with examples of various sceptical arguments from ignorance. These arguments are at the heart of the problem of sceptical paradoxes, and providing a satisfactory account of how they work is an important part of seeking happy-face resolution to this problem. This chapter focuses on the following question: what epistemic principle underwrites sceptical arguments from ignorance? In chapter 2, I was concerned with give a general account of the effective sceptical hypotheses that feature in sceptical arguments from ignorance. So too, in this chapter, I am concerned with the prospect of providing a general account of the epistemic principle underwriting sceptical arguments from ignorance.

I begin, in section 3.2, by setting out some preliminary points clarifying what is involved in the general account of sceptical arguments sought here. Next, in section 3.3, I consider some existing proposals for the principles underwriting sceptical arguments from ignorance: *Closure*, *Underdetermination*, and *Infallibility*. I consider each in turn, looking at how the principle applies to some paradigm sceptical arguments. In section 3.4, I go on to critically assess the adequacy of these principles as accounts of sceptical arguments from ignorance, and consider another proposed principle, INFALLIBILITY+, that avoids these adequacy issues. Finally, in section 3.5, I raise an issue with *Infallibility*+, and forward my own proposal for the principle underwriting sceptical arguments from ignorance. I argue that my principle is preferable to the alternatives by showing

how it applies to various examples, and addresses the explanatory issue I raised.

3.2 Preliminaries.

In what sense will an account of the epistemic principle underwriting sceptical arguments from ignorance count as a general one?¹ Both straightforward challenges to specific knowledge claims (e.g. BICYCLE) and those sceptical challenges (e.g. ENVATTED BRAIN and DREAMING) which target virtually all knowledge claims at once can be expressed in terms of AI. The latter sort of arguments give rise to paradoxes, and so, I think it is desirable to seek a general account of these arguments' common core to develop our understanding of how to solve these paradoxes.

Compare these arguments with the following argument:²

ZOO.

(7) S does not know that the animals in the pen are not cleverly disguised mules.

(8) If S knows that this animal in the pen is a zebra, then S knows that the animals in the pen are not cleverly disguised mules.

So,

(9) S does not know that this animal in the pen is a zebra.

ZOO is similar to ENVATTED BRAIN, and differs from DREAMING in that the sceptical hypothesis involved (that *the animals in the pen are cleverly disguised mules*)—as I noted in the previous chapter—entails that the targeted proposition (i.e. this animal in the pen is a zebra) is false. Nonetheless, as we have seen, the sceptical hypothesis in the ZOO argument can be changed so as to *not* entail the falsity of the target proposition. The hypothesis that *half of the animals in the pen are cleverly disguised mules*, for example, does not entail that *this animal in the pen is a zebra*. Moreover, ZOO differs from ENVATTED BRAIN and DREAMING in two important ways. The first of these concerns the scope of the sceptical challenge involved. This, as has been noted, is narrower than the latter two sceptical arguments; ZOO does not threaten to challenge all our

¹ Compare (Wright, 2005), where a distinction between 'Cartesian' and 'Humean' sceptical arguments is made. I am here concerned only with the former.

² This argument is based on considerations drawn from (Dretske, 1970). Similar considerations feature predominantly in the discussions of scepticism. See, for example, (Wright, 2002).

knowledge claims at once. The second way in which ZOO differs is that, unlike the latter two examples of sceptical challenges, this argument is not necessarily paradoxical. To illustrate this, suppose S is a geneticist who is standing at the zebra pen having just completed an on-site DNA analysis of all the zoo's mammals. Suppose further that S is aware that the result of analysing the DNA from the animal that they are looking at means that the animal is not a mule. It seems now that the first premise of ZOO lacks plausibility, and the argument, subsequently, does not present us with a paradox. Importantly, it seems that in this case there is a satisfactory way to counter the argument. In this case, for example, it seems perfectly legitimate to respond to ZOO by saying "But S knows that this animal is not a mule!".

I think that a satisfactory account of the epistemic principles underwriting sceptical arguments such as ENVATTED BRAIN and DREAMING should also apply to examples such as ZOO. In the latter, but not the former cases, these arguments can be potentially countered, and are not paradoxical. Moreover, some AI arguments can be used to challenge a claim that 'S knows that p ' irrespective of whether sh entails the falsity of p (e.g. ENVATTED BRAIN, ZOO) or not (e.g. DREAMING). I propose that any adequate account of the principles underlying sceptical arguments from ignorance should at least accommodate, if not explain, each of these facts. In the next section, I outline some candidates for the epistemic principles underwriting the sceptical arguments from the previous section, and briefly comment on how they might apply to ENVATTED BRAIN, DREAMING, and ZOO.

3.3 Candidate epistemic principles.

3.3.1 Closure.

The first candidate for the epistemic principle underwriting sceptical arguments can be called *epistemic closure*. Simply put, the idea behind this principle is this. Someone knows those propositions that they come to believe on the basis of deducing them from other propositions that they know. We can express the principle more formally in the following way:

Closure

For all S, p , q , if S knows that p and competently deduces q from p , thereby coming

to believe q , while retaining their knowledge that p , then S comes to know that q .³

The principle basically says that whenever someone knows some proposition which entails another proposition, and this person forms a belief in the latter because they recognised that it was entailed by the former, then that person also knows the latter so long as they continue to know the former. So, what bearing does *Closure* have on explaining how arguments of the AI form really work? To see how, we now look at how this principle is thought to explain two of our sceptical argument examples from above; ENVATTED BRAIN and ZOO. First, it is noted that the first premises just seem to be plausible. This can be supported by an appeal to what appear to be relatively uncontroversial assumptions. One example would be the following claim: no one is able to tell whether they are a brain in a vat or not. Hence the plausibility of ENVATTED BRAIN's first premise. Likewise, it just seems uncontroversial that most people would be unable to tell the difference between a cleverly disguised mule and a zebra. Hence the plausibility of ZOO's first premise.

Both argument's second premises can then be motivated by appeal to *Closure*. To do this, we start by noting that we can come by means of rational reflection to grasp the following relevant entailments. First, that *S has hands* entails that *S is not a handless brain in a vat*. Second, that *this animal in the pen is a zebra* entails that *the animals in the pen are not cleverly disguised mules*. Next, we note that through rational reflection we could come to believe the entailed proposition on the basis of deducing it from, and without ceasing to know (if they already did), the former proposition. Now, *Closure* tells us that in every case where S knows that they have hands, S also knows that they are not a handless brain in a vat. And whenever S knows that this animal in the pen is a zebra, they also know that the animals in the pen are not cleverly disguised mules. Hence the plausibility of the second premises of ENVATTED BRAIN and ZOO.

Does the *Closure* principle provide an adequate account of how sceptical arguments from ignorance work? I think that it does not. I think that it fails in this respect for the reason that it fails to apply to a large proportion of sceptical arguments from ignorance. I explain why later on, but first I briefly outline two other candidates for the epistemic principle underwriting sceptical arguments from ignorance.

³ Cf. (Almeida, 2012, 199), (Brueckner, 2005, 388), (Brueckner, 1994, 827), and (Cohen, 1998, 144). For formulations in terms of justification see (Heller, 1999, 196), (McCain, 2013, 291), and (Pritchard, 2005, 107). For a comprehensive discussions of epistemic closure see (Hawthorne, 2004).

3.3.2 Underdetermination.

Our second candidate for the epistemic principle underwriting sceptical arguments can be called *epistemic underdetermination*. Put simply, the idea behind this principle goes something like this: someone knows a proposition p only if the basis on which they believe that proposition tells against the truth of propositions whose falsity is entailed by p . We can express the principle more formally in the following way:

Underdetermination

For all S, p, q , if S 's evidence for believing p does not favour p over some q which S knows to be incompatible with p , then S 's evidence does not justify S in believing p .⁴

The principle basically says that whenever someone knows some proposition then they possess evidence which favours this proposition over any proposition which entails that p is false. So, what bearing does *Underdetermination* have on explaining how arguments of the AI form really work? To see how, we will do what we did with the previous candidate principle, and examine how *Underdetermination* is thought to explain ENVATTED BRAIN and ZOO.

Let us take ENVATTED BRAIN first. Its first premise may be motivated by the following instance of *Underdetermination*, and a relatively uncontroversial sceptical assumption:

(U1) If S 's evidence for believing that S is not a brain in a vat does not favour that proposition over the proposition that S is a brain in a vat, then S 's evidence does not justify S in believing that S is not a brain in a vat.

(U2) No one's evidence for believing that they are not a brain in a vat favours the proposition that *they are not a brain in a vat* over the proposition that *they are a brain in a vat*.

From (U1) and (U2) it follows that no one has evidence that justifies them in believing that they are not a brain in a vat. Given only the relatively uncontroversial assumption that knowing a proposition requires having evidence that justifies you in believing that proposition, it follows that no one knows that they are not a brain in a vat. In this way it is suggested we can explain the first premise of ENVATTED BRAIN; that S does not *know*

⁴ Adapted from (Boult, 2013, 1127). See also (Cohen, 1998, 144-5), (Pritchard, 2005, 108), and (Williamson, 2000).

that they are not a brain in a vat. Similar considerations apply to ZOO; the plausibility of ZOO's first premise falls out of the following instance of *Underdetermination* and a relatively uncontroversial sceptical assumption:

(U3) If S's evidence for believing that \neg *disguised mule* does not favour the proposition that \neg *disguised mule* over the proposition that *disguised mule*, then S's evidence does not justify S in believing that \neg *disguised mule*.

(U4) Practically no person's evidence for believing that *the animal in the pen is not a cleverly disguised mule* favours that proposition over the proposition that *the animal in the pen is a cleverly disguised mule*.

The second premise of ENVATTED BRAIN may be motivated in terms of the same principle. First, we note that, the proposition that *S is a brain in a vat* entails that it is false that *S has hands*. So, the two propositions are incompatible. Next, we note that, this incompatibility is knowable by virtually everyone. And so, we can see that the antecedent of ENVATTED BRAIN's second premise will be true just in case S's evidence for believing that *S has hands* favours that proposition over the proposition that *S is a brain in a vat*. In which case S would be in a position to know that they are not a brain in a vat. After all, their true belief that they are not will be justified in virtue of their evidence which favours it over its negation. Hence the plausibility of ENVATTED BRAIN's second premise. Similar considerations apply to ZOO; the incompatibility of the hypothesis that *the animals in the pen are cleverly disguised mules*, and the proposition that *this animal in the pen is a zebra* is easily knowable through rational reflection by anyone. According to the underdetermination principle, S can have evidence justifying their belief that the animal is a zebra only if they have evidence for *this* proposition that favours it over the proposition that *the animals in the pen are cleverly disguised mules*. Hence the plausibility ZOO's second premise.

Does the *Underdetermination* principle provide an adequate account of how sceptical arguments from ignorance work? I think that it does not, and that it fails in this respect for the same reason as *Closure*. We will see why shortly, but first let's briefly look at one final candidate for the epistemic principle underwriting sceptical arguments from ignorance.

3.3.3 Infallibility.

Our third candidate for the epistemic principle underwriting sceptical arguments can be called *epistemic infallibility*. Put simply, the idea behind this principle goes something like this. Someone knows a proposition only if they have eliminated all those propositions whose falsity is entailed by p . We can express the principle more formally in the following way:

Infallibility.

For all S, p, q , S knows that p only if S 's evidence eliminates every possibility in which p is false.⁵

The principle says that whenever someone knows some proposition then their evidence eliminates every possibility in which this proposition is false. To illustrate, the infallibility principle holds that if I know that I have hands, for example, then my evidence eliminates every possibility in which I do not have hands. So, what bearing does *Infallibility* have on our understanding of how arguments of the AI form really work? To explore this question, let's again follow the same procedure as we did with *Closure* and *Underdetermination*. Here, then, we want to apply the *Infallibility* principle to the examples of ENVATTED BRAIN and ZOO. What we are interested in exploring is how, if at all, the candidate epistemic principle serves to motivate the various premises of AI arguments.

Starting again with ENVATTED BRAIN. Its first premise is motivated by the following instance of *Infallibility* and a relatively uncontroversial sceptical assumption:

(I1) S knows that they are not a brain in a vat only if S 's evidence eliminates every possibility in which the proposition that *S is not a brain in a vat* is false.

(I2) No one can eliminate the possibility that they are a brain in a vat.

The possibility in which S is a brain in a vat is one in which the proposition that *S is not a brain in a vat* is false. So, from (I1) and (I2), it follows that S does not know that they are not a brain in a vat. Hence the plausibility of ENVATTED BRAIN's first premise. And similar considerations apply to ZOO's first premise. The possibility in which the animal in front of S at the zoo is cleverly disguised mule is one in which it is false that *the animal in the pen is not a cleverly disguised mule*. So, in order to know that the animal

⁵ Adapted from (Lewis, 1996, 551).

in the pen is not a disguised mule, S must eliminate the possibility that proposition is false. Yet it is a further relatively uncontroversial assumption that most people in S's position at the zoo would be unable to eliminate this possibility, hence the plausibility of ZOO's first premise.

The second premise of ENVATTED BRAIN can also be seen to be motivated by the *Infallibility* principle. To see how, we need to first note that the proposition *S has hands* will be *false* whenever S is a brain in a vat. And so, the antecedent of ENVATTED BRAIN's second premise will be true just in case S has eliminated that *S is a brain in a vat*. In which case, S would be in a position to know that they are not a brain in a vat. After all, they have eliminated the very possibility of it being false. Hence the plausibility of ENVATTED BRAIN's second premise. Similar considerations apply to ZOO. Any possibility in which the animal that S thinks is a zebra is in fact a cleverly disguised mule is one in which it is false that *the animal S thinks is a zebra is, in fact, a zebra*. So, according to *Infallibility*, the antecedent of ZOO's second premise will be true just in case S has eliminated the proposition that *the animal in the pen is a cleverly disguised mule*. In which case, S will be in a position to know that the animals in the pen is not a cleverly disguised mule; they'll have eliminated every possibility in which it is false. Hence the plausibility ZOO's second premise.

Does the *Infallibility* principle provide an adequate account of how sceptical arguments from ignorance work? I think that it does not, and that it also fails in this respect for the same reason as *Closure* and *Underdetermination*. I discuss this failure next.

3.4 Principle of Ignorance Possibilities.

In this section, I critically assess the relative merits of *Closure*, *Underdetermination*, and *Infallibility* as candidates for the epistemic principle underwriting the core sceptical argument form AI.

The case against the principles considered above as accounts of those underwriting the premises of AI can be stated straight-forwardly. First, an adequate account of the principle underwriting AI should apply to each of the instances presented in the previous section. Second, neither *Closure*, *Underdetermination*, nor *Infallibility* apply to the second premise of the DREAMING argument. To see why, let *p* be the proposition that *S has hands*. The second premise of the DREAMING argument, then, says that if S knows that they have hands, then they know that they are not currently dreaming. Neither *Closure*, *Underdetermination*, nor *Infallibility* can underwrite this premise, as

they do the corresponding premise of ENVATTED BRAIN. This is because the hypothesis that *S* is dreaming is not incompatible with the proposition that *S* has hands. After all, anyone can be currently dreaming and have hands. I assume that an adequate account of the principle underwriting sceptical arguments from ignorance should apply to as many recognised instances as possible. So, on my view, neither *Closure*, *Underdetermination*, nor *Infallibility* provides an adequately general account of the epistemic principle underwriting sceptical arguments from ignorance.

The argument is simple, and has support in recent literature on the topic. For example, Kraft (2013) argues in a similar way against these principles, and uses this to motivate an alternative. For Kraft it is clear that AI cannot be adequately explained in terms which only apply to those instances where the sceptical hypothesis entails that *p* is false. This is evident from the following comment:

[T]he most charitable interpretation of sceptical scenarios understands them as illustrating the possibility of global (empirical) ignorance, the possibility of not knowing anything... A sceptical scenario is not defective just because it does not illustrate the possibility of global error. To be successful it has to illustrate the possibility of global (empirical) ignorance, i.e. the possibility that all of my (empirical) beliefs fall short of knowledge' (Kraft, 2013, 66)

Before commenting further, it is worth noting that Kraft takes a more restrictive view of what counts as a sceptical hypothesis than I do. On my view, as outlined in chapter 2, sceptical hypotheses are those that illustrate the possibility of ignorance with respect to a *range* of propositions. I think that Kraft's comments, however, are informative in the sense that he emphasises that sceptical hypotheses need not be possibilities of error. The idea, here, is that a possibility of error is one *in which* a proposition that is claimed to be known *is false*; in other words, a possibility that is *incompatible* with a putatively known proposition *p*. In this way, I take it that Kraft agrees with my argument for the inadequacy of those accounts of sceptical arguments in terms of epistemic principles which do not apply to examples where the sceptical hypothesis and targeted proposition *p* are compatible. Accounts in terms of *Closure*, *Underdetermination*, and *Infallibility* are all found wanting on this view.

Kraft, however, thinks that an adequate account of the epistemic principle underwriting sceptical arguments can be given. His strategy for doing so is simply to adapt the infallibility principle to accommodate those sceptical arguments to which *Infallibility*

failed to apply. Kraft's approach is to substitute *Infallibility* with what he terms "the principle of excluded ignorance-possibilities. According to this principle, S knows that P only if S can rule out all possibilities in which S does not know that P" (Kraft, 2013, 67). In support of this principle, Kraft points out that the same considerations that appear to motivate *Infallibility* also motivate the principle of excluded ignorance-possibilities. Kraft notes that the *Infallibility* principle, or what he calls the "the principle of excluded error-possibilities [which says]... one knows that P only if one can eliminate all possibilities in which P is false" (Kraft, 2013, nt. 14) is well motivated. He explains that:

Lewis defends the principle of excluded error-possibilities by pointing out that a concessive knowledge attribution like "S knows that P, but she has not ruled out the possibility that P is false" "just sounds contradictory". (Kraft, 2013, 68)

And the same considerations also motivate 'the principle of excluded ignorance-possibilities', or what I term *Infallibility+*:

Infallibility+.

For all S, *p*, S knows that *p* only if S eliminates every possibility in which S does not know that *p*.

For example, each of the following claims sound no less strange and contradictory; 'S knows that they have hands, but does not know that they are not currently dreaming', 'S knows that the Bank is open, but does not know that the Bank has not changed its opening hours since they last checked'. So, the considerations that motivated *Infallibility* also seem to motivate *Infallibility+*. In fact, as Kraft notes, the latter principle implies the former:

On my account, ignorance is lack of knowledge, i.e. S is ignorant that P if it is not the case that S knows that P. A useful side-effect of this stipulation is that on my account all error possibilities are *ipso facto* ignorance-possibilities (but not *vice versa*). (Kraft, 2013, nt. 13)

The point here is significant. This is because it seems that *Infallibility+* does not face the problems which I argued beset *Closure*, *Underdetermination*, and *Infallibility*. This was the problem of failing to account for sceptical arguments from ignorance *generally*, in the sense that they failed to apply to examples, such as DREAMING, in which the sceptical

hypothesis does not entail the falsity of the targeted proposition. In the remainder of the this section, I will set out how *Infallibility+* applies generally to the various examples of sceptical argument we are considering. In the next section, I move on to suggest a potential reason to be dissatisfied with *Infallibility+*, and draw on this point to propose my own alterantive account of the epistemic principle underwriting sceptical arguments from ignorance.

How does *Infallibility+* apply to the examples of sceptical arguments that we have been considering? The principle says that whenever someone knows some proposition then they have eliminated every possibility in which they do not *know* this proposition. To illustrate, *Infallibility+* holds that if I know that I have hands, for example, then I have eliminated every possibility in which I do not *know* that I have hands. Let's follow the same basis procedure as we did with *Closure*, *Underdetermination*, and INFALLIBILITY. I will, again, apply *Infallibility+* to the examples of ENVATTED BRAIN, DREAMING, and ZOO.

Starting with ENVATTED BRAIN. Its first premise is motivated by the following instance of *Infallibility+* and a relatively uncontroversial sceptical assumption:

(I+1) S knows that they are not a brain in a vat only if S eliminates every possibility in which S does not know that they are not a brain in a vat.

(I+2) No one can eliminate the possibility that they are a brain in a vat.

The possibility in which S is a brain in a vat is one in which S does not know that *they are not a brain in a vat*—because that very proposition is false. So, from (I+1) and (I+2), it follows that S does not know that they are not a brain in a vat. Hence the plausibility of ENVATTED BRAIN's first premise. Similar considerations apply to ZOO's first premise. The possibility in which the animal in front of S at the zoo is trivially one in which it is false that *the animal in the pen is not a cleverly disguised mule*. Consequently, this possibility is trivially one in which S does not *know* that *the animal in the pen is not a cleverly disguised mule*. So, in order to know that the animal in the pen is not a disguised mule, S must eliminate the possibility that the animal in the pen is a cleverly disguised mule. The plausibility of ZOO's first premise—that S does not know that the animal in the pen is not a cleverly disguised mule can be accounted for, then, by *Infallibility+* and a sceptical assumption corresponding to (I+2)—that most people in S's position at the zoo would not eliminate this possibility.

The second premise of ENVATTED BRAIN can also be seen to be motivated by the *Infallibility+* principle. To see how, we need to first note that S will not know that *S has*

hands whenever S is a brain in a vat.⁶ And so, the antecedent of ENVATTED BRAIN's second premise will be true just in case S has eliminated that *S is a brain in a vat*. In which case, presumably, S would also be in a position to eliminate any possibility in which they do not know that they are not a brain in a vat on the basis of having eliminated know that they are not a brain in a vat—after all, they have eliminated *any* possibility in which they are one. And this could, then, help to account for the plausibility of ENVATTED BRAIN's second premise. Similar considerations apply to ZOO's second premise. Any possibility in which the animal that S thinks is a zebra is, in fact, a cleverly disguised mule is one in which S does not know that the animal is a zebra. So, according to *Infallibility+*, the antecedent of ZOO's second premise will be true just in case S has eliminated the proposition that *the animal in the pen is a cleverly disguised mule*. In which case S will be in a position to know that the animal in the pen is not a cleverly disguised mule. They will have eliminated every possibility in which it is false, and thereby be able to eliminate any possibility in which they do not know that the animal in the pen is not a cleverly disguised mule.

So far, then, it seems that *Infallibility+* fares no worse than either *Closure*, *Underdetermination*, or *Infallibility* in terms of providing some account of the premises of sceptical arguments from ignorance such as ENVATTED BRAIN and ZOO. I suggested above that it is an advantage of *Infallibility+* over those alternative principles that it also provides a more general account of sceptical arguments in virtue of applying to examples that involve sceptical hypotheses that do not entail the falsity of the targeted proposition *p*. This can now be seen by applying this principle to the example of DREAMING.

With respect to DREAMING's first premise, *Infallibility+* states that S knows that S is not dreaming only if they have eliminated the corresponding ignorance possibility that they do not know that they are *not* dreaming. The premise that S does not know that they are not dreaming can be motivated, then, given only the further sceptical assumption that no one is in a position to eliminate the possibility that they are dreaming. This is because any possibility in which S is dreaming is *also* one in which S does not know that they are *not* dreaming. If S has not eliminated that they are dreaming, then they have not eliminated the possibility that they do not know that they are not dreaming, and so, consequently, by *Infallibility+*, S will not know that they are not dreaming. And similar considerations to those discussed seem to help account for the second premise of DREAMING. We begin by noting that any possibility in which the S is dreaming is one in which S does not know that they have hands. So, according to *Infallibility+*, the

⁶ On my view, this is because nothing is known by any S who is a brain a vat. Cf. chapter 2.

antecedent of DREAMING's second premise will be true just in case S has eliminated the possibility that they are dreaming. In which case, S will be in a position to know that they are not dreaming for similar reasons to those I noted above. They will be in a position to eliminate every possibility in which they do not know that they are not dreaming on the basis of having eliminated any possibility that they *are* dreaming.

It seems, then, that *Infallibility*⁺ can provide an account of sceptical argument in a way that the other candidate principles could not. It does so in a way that applies regardless of whether *sh* entails that a targeted proposition *p* is false or not. I think, however, that there is a distinct issue facing an attempt to account for sceptical arguments from ignorance in terms of *Infallibility*⁺, which I discuss in the next section.

3.5 Principle of Epistemic Possibility.

In this section, I suggest some issues surrounding an *Infallibility*⁺ based account of the epistemic principle underwriting sceptical arguments from ignorance, and use this critique to motivate my own proposal in terms of a novel analysis of *epistemic possibility*. First, I outline what I take to be an issue surrounding the explanatory adequacy of this account by drawing on a critical point from DeRose that was already noted in chapter 2. Next, I use this issue as a springboard to propose an alternative epistemic principle based on an analysis of epistemic possibility. I argue that my proposal represents a more attractive account than its competitors of the principle underwriting sceptical arguments from ignorance, by showing how it both applies generally to the examples we have been considering, and addresses the explanatory issue facing *Infallibility*⁺.

Consider again, as we did in section 2.2, the following critical insight from DeRose:

[A] treatment of AI must tell us what it is about skeptical hypotheses that makes it difficult to claim to know that they don't obtain. The key feature of skeptical hypotheses... is clearly this: we can't rule them out... It is indeed plausible to suppose that we can't rule out skeptical hypotheses. And it's plausible that we don't know that they don't obtain. But it doesn't seem to advance our understanding much to explain the plausibility of either by that of the other... we need an explanation that reaches outside this circle of all-too-closely related terms of epistemic appraisal. (DeRose, 1995, 16-7)

DeRose's point here is, I think, significant. All things being equal, it would be preferable to find a way of explicating the notion of eliminating a possibility which is informative in

the sense that DeRose describes. That is, it would be more informative to explain what we mean when we say that S has not eliminated a possibility which goes beyond simply saying that S does not know that possibility to be false. Consider the following comment:

One natural way of understanding what is involved in ‘ruling out’ an error-possibility is that of the agent knowing that error-possibility to be false. Indeed, we often treat these two notions as equivalent. (Pritchard, 2005, 25)

Whilst Pritchard’s comment, here, explicitly mentions *error*-possibilities, I think we can develop the point to apply to ignorance-possibilities. That is, a natural way of understanding what is involved in S eliminating or ‘ruling out’ an ignorance-possibility, is that of S knowing that ignorance-possibility to be false. DeRose’s point in the preceding quoted passage is, I take it, that treating these notions as equivalent is uninformative as far as explaining the *principled* plausibility of the sceptical arguments’ first premises. And, all things being equal, we should prefer that our explanation is more informative than *that*. I think that we can give an explanation that is more informative in this sense, and that it can be found in an analysis of the notion of epistemic possibility. Before setting out my proposal, it remains to briefly run through the reasons why I think that considerations of the DeRose and Pritchard points present an issue with the *Infallibility*+ account of sceptical arguments.

Suppose, given our interest in understanding the premises of sceptical arguments we ask, for example, ‘Why is it that I do not know that I am not a brain in a vat?’ If the answer were that ‘because you cannot eliminate the possibility that you are a brain in a vat’, I think we would have reason to be dissatisfied by this explanation. And this is because we could, then, naturally inquire further and ask ‘but why am I not able to eliminate that possibility?’. The point to be drawn from the comments of DeRose and Pritchard, here, is that knowing something is normally understood as equivalent to eliminating that it is false. As such, the standard answer to this second question might well be ‘because you do not know that it does not obtain’. And *that*, it seems, does not help us understand very much at all. Moreover, I think that similar considerations will lead to our dissatisfaction with an account of sceptical arguments in terms of *Infallibility*+. To illustrate this, let us imagine what sort of answers to the proceeding line of inquiry that we could expect from an *Infallibility*+ account. Consider, then, the following:

Q1. Why don’t I know that I am not a brain in a vat?

A1. Because you haven’t eliminated that you don’t know that you are a brain in vat.

Q2. Why haven't I eliminated that I don't know that I am not a brain in a vat?

A2. Because you don't *know* that you haven't eliminated that you don't know that you are one.

On the standard understanding of what is involved in 'eliminating' a possibility, noted above, I take it that the answer A2 is a legitimate response to Q2. But it is not, to my mind, a particularly explanatory, and subsequently, satisfactory answer to the question. And the situation becomes worse, in this respect, if we seek to understand and naturally inquire further. To illustrate, consider the following development of the preceding line of questioning:

Q3. Why don't I *know* that I haven't eliminated that I don't know that I am one?

A3. Because you haven't eliminated *that* you don't know that you haven't eliminated that you don't know that you are a brain in vat.

Here, again, the answer A3 would seem to be a legitimate response to Q3, licenced by an analysis that treats 'eliminating' a possibility as equivalent to knowing it to be false. The *Infallibility+* account explains our lack of knowledge in terms of our failing to eliminate the corresponding ignorance possibilities. But if eliminating is treated as equivalent to knowing to be false, it seems that our natural line of inquiry will only take us further from any satisfactory answer or understanding. If A3 was not already obscure enough, it does not take long before our natural questions lead to bizarre and confusing iterations:

Q4. Why haven't I eliminated *that* I don't know that I haven't eliminated that I don't know that I am not a brain in vat?

A4. Because you don't *know* that you haven't eliminated that you don't know that you haven't eliminated... etc.

A vicious regress, of sorts, seems threatens so long as the questions remain, and the answers are given in terms of *Infallibility+* and the standard analysis of eliminating a possibility. At very least, I do not think that any particularly satisfying level of understanding of sceptical premises will be forthcoming. I do not intend to argue that these considerations represent an insurmountable objection to an *Infallibility+* account.

Rather, I think an alternative account can be had that addresses these concerns, and which applies no less generally to sceptical arguments from ignorance.

Drawing on the preceding discussion, I now set out my preferred account of the principle underwriting sceptical arguments from ignorance. I propose the following as a plausible constraint on knowledge:

Epistemic Possibility.

For all S, p , S knows that p only if it is not epistemically possible for S that $\neg p$.

The principle says that whenever someone knows some proposition then the falsity of that proposition is not epistemically possible for that person. To illustrate, *Epistemic Possibility* holds that if I know that I have hands, for example, then it is not the case that it is epistemically possible for me that *I do not have hands*. Below, I will run through how *Epistemic Possibility* applies to the examples of ENVATTED BRAIN, DREAMING, and ZOO. Before doing this, however, it is important to comment on what analysis of epistemic possibility I am adopting here. As I will explain shortly, this has a significant bearing on how I think we can address the sort of explanatory concern I have attributed to DeRose above. There does not appear to be an universally accepted analysis of epistemic possibility, but, for my purposes here, I adopt the following analysis:⁷

A proposition p is epistemically possible for a subject S iff p is consistent with S knowing what they do.

In this way, the *Epistemic Possibility* principle can be understood as saying that whenever someone knows a proposition, then that proposition cannot be false *given* the facts about what that person knows.⁸ With this analysis in place, let's look at how *Epistemic Possibility* applies to our various examples of sceptical arguments. The first premise of ENVATTED BRAIN could be motivated by the following instance of this principle, and a corresponding sceptical assumption:

(EP1) S knows that they are not a brain in a vat only if it is not epistemically possible for S that they are a brain in a vat.

⁷ For various interesting discussions of epistemic possibility, see (Anderson, 2014), (Dougherty and Rysiew, 2009), (Huemer, 2007), (Przyjowski, 2017), (Reed, 2013), and (Wright, 2007).

⁸ I do attempt to provide an in-depth defence of this analysis here. For my present purposes, I think it is enough to note that this is not obviously wrong. I focus, instead, on how I think this analysis can help us to better understanding how sceptical arguments from ignorance work.

(EP2) For anyone, it is epistemically possible that they are a brain in a vat.

The plausibility of ENVATTED BRAIN's first premise—that S does not know that they are not a brain in a vat—is motivated straight-forwardly on this account, given only (EP1) and the sceptical assumption (EP2). Moreover, this assumption itself is, I think, representative of the natural things we might say about our epistemic position with respect to various sceptical hypotheses: that, *for all we know*, they do obtain.⁹ Likewise for the first premises of both ZOO and DREAMING. In either case, I take it that a natural way of explaining why I might not know the relevant sceptical hypothesis to be false will involve pointing out that *for all I know*, these hypotheses are *true*, or, that they are *epistemically possible for me*. On my analysis of epistemic possibility, this amounts to saying that it is consistent with me knowing everything I *do* that *I am currently dreaming*, for example, or that *the animal in the pen is a cleverly disguised mule*.

The second premises of each sceptical argument we have been considering can also be seen to be motivated by considerations involved in the *Epistemic Possibility* principle. The reasons for this are related, as I shall explain, to the considerations of sceptical hypotheses set out in chapter 2. On my view, every effective sceptical hypothesis describes a possibility *in which* the falsity of some targeted empirical proposition *p* is *consistent with* S knowing what they do in that scenario. This is trivially true in the case of the *dreaming* and *brain in a vat* hypotheses; if either were true of S, then S would not know any empirical proposition at all. Consequently, the *brain in a vat* and *dreaming* hypotheses describe scenarios *in which* it is consistent with S knowing what they do, that $\neg p$, for *any* empirical proposition *p*—viz. every empirical proposition *p* is *epistemically possible* for S.

How do these points relate to how the second premises of sceptical arguments are underwritten by *Epistemic Possibility*? My central idea here is, put coarsely, that the *Epistemic Possibility* principle, on my analysis, can be brought together with the insight that sceptical hypotheses are possibilities of epistemic possibility. And I think that doing so can help to develop a more explanatory account of how the major premises of sceptical arguments are underwritten. In order to explicate this, I think it is instructive to consider the point that “anything compatible with something that is compatible with *p* is compatible with something that leaves the *possibility* of *p* open” (Brandom, 2008, 129). This point bears importantly on the analysis of epistemic possibility that I am adopting, and brings into view an idea about how our epistemic positions with respect to various propositions and corresponding sceptical hypotheses are related. The idea, here,

⁹ Cf. (DeRose, 1995, 2).

is that the following seems *prima facie* plausible: if $\neg p$ is inconsistent with S knowing what they do, and q entails that $\neg p$ is consistent with S knowing what they do, then q is inconsistent with S knowing what they do. Translating this into terms of epistemic possibility, we get the following:

If it is not epistemically possible for S that $\neg p$, and q entails that it is epistemically possible for S that $\neg p$, then it is not epistemically possible for S that q .

What bearing does this have on explaining how the major premises of sceptical arguments are underwritten? The *brain in a vat* hypothesis entails that it is epistemically possible for me that I do not have hands. As such, if it is not epistemically possible for me that I do not have hands, then it is not epistemically possible for me that the *brain in a vat* hypothesis is true. Now, *Epistemic possibility* says that I know that I have hands only if it is not epistemically possible for me that I do not have hands. So, if I know I have hands, then it is not epistemically possible for me that the *brain in a vat* hypothesis is true. Given only the further assumption that we will be prepared to say I know a proposition if I believe it, and it is not epistemically possible for me that it is false, I take it that *Epistemic Possibility* tells us that if I know I have hands, then I am in a position to know that I am not a brain in a vat. The same sort of explanation applies to the second premise of DREAMING; the *dreaming* hypothesis entails that every empirical proposition is epistemically possible for S, and so, given the same considerations just outlined, it seems that S is in a position to know that the *dreaming* hypothesis is false if they know any given empirical proposition.

What about the second premise of ZOO? We can recognise that hypothesis that *the animal in the pen is a cleverly disguised mule* entails that it is epistemically possible for S that *the animal in the pen is not a zebra*. We can also recognise that if it is not epistemically possible for S that the animal in the pen is not a zebra, then it is also not epistemically possible for S that the animal in the pen is a cleverly disguised mule. Again, *Epistemic Possibility* holds that S knows that the animal in the pen is a zebra only if it is not epistemically possible for S that the animal in the pen is not a zebra. So, if S knows that the animal is a zebra, then it is not epistemically possible for S that the *cleverly disguised mule* hypothesis is true. On the assumption that we will be prepared to say someone knows a proposition if they believe it, and it is not epistemically possible for them that it is false, I take it that this helps to explain our inclination to judge that if S knows that the animal in the pen is a zebra, then S also knows that the animal is not a cleverly disguised mule.

It remains to set out how I think my proposed principle *Epistemic Possibility* addresses the explanatory concerns I raised above for the *Infallibility+* account. On my view, there is no obvious sense in which my account of how sceptical premises are underwritten is beset by the sort of explanatory circularity or regress highlighted above. To the contrary, I think that my proposed principle and my analysis of epistemic possibility together can help to provide us with a deeper understanding of sceptical arguments. To illustrate this, consider, then, the following:

Q5. Why don't I know that I am not a brain in a vat?

A5. Because it is epistemically possible for you that you are a brain in vat.

Q6. Why is it epistemically possible for me that I am a brain in a vat?

A6. Because it is consistent with you knowing what you do, that you are a brain in a vat.

On my view, A5 and A6 are legitimate answers to questions Q5 and Q6, but are they satisfactorily explanatory? There are a couple of things to say here. First, I am not convinced that there are any further obvious questions to ask in response to A6. The answer does not seem to require further explanation, at least not in terms of the notions involved. We can naturally understand what A6 says at face value. It says, ultimately, that there is no empirical proposition that I know, which I would not *if* I was a brain in a vat. And that is tantamount to saying that I do not know any empirical proposition at all.¹⁰ Now, this explanation of sceptical premises might seem to beg the question, and perhaps, consequently, it might prompt us to ask for further explanation. I address this shortly in my second point, but before I do I want to note that there is an important sense in which my analysis of epistemic possibility allows us to 'break out of the circle of all too closely related terms'. It does so in virtue of understanding epistemic possibility as a function of our knowing what we do; we can assess claims of epistemic possibility easily, provided we can identify what we do and don't know, and recognise what can and cannot, subsequently, be the case. This leads on to my second point. A6 might appear to be unsatisfactorily explanatory in the context of understanding sceptical arguments; it may strike us as question begging. We might still want to inquire further in response to A6, and the following is what I take to be a natural example of how we might do this:

¹⁰ See section 6.8. Compare this Williamson's reconstruction of the sceptic's reasoning, and his response to it, in chapter 8 of (Williamson, 2000). In particular, his observation that the claim of ID discriminability "is tantamount to the sceptic's conclusion. The sceptic cannot use it as a premise without begging the question" (Williamson, 2000, 168).

Q7. Why is *that* though; why don't I know, for example, that I have hands?

A7. Because it is consistent with you knowing what you do that you don't have hands.

Again, on my view, A7 is a legitimate response to Q7, but is it satisfactorily explanatory in terms of how sceptical premises are underwritten? For my part, I think both that it is, and that it can help to illuminate the perplexing nature of sceptical arguments. To see this, it is worth noting that in responding to A6, we could naturally substitute any empirical proposition *q* for the proposition that *I have hands* to the same effect. Importantly, on my view, an answer corresponding to A7 involving *q* will be a legitimate response on behalf of the sceptic. In the final analysis, the sceptic seems to rely on the claim that for any empirical proposition *q* epistemically possible for you that *q* is false, or in other words, that *we know no empirical propositions at all*. Is this satisfactorily explanatory? I conclude this section by noting that, for my purposes, my account of the epistemic principle underwriting sceptical arguments is not obviously *less* explanatory than *Infallibility+* in this respect. Indeed, I aim to show in chapter 6 that it can contribute to an understanding of how to satisfactorily resolve sceptical paradoxes.

3.6 Concluding remarks.

In this chapter, I outlined and briefly commented on various examples of sceptical arguments that feature prominently in the literature. I began, in section 3.2, by suggesting that a general account of sceptical arguments is desirable, and setting out what such an should apply to ENVATTED BRAIN, DREAMING, and ZOO. In sections 3.3 and 3.4, I critically examined *Closure*, *Underdetermination*, and *Infallibility*, and considered how the alternative *Infallibility+* performed better with respect to my desiderata. Finally, in section 3.5, I argued that *Infallibility+* faced a potential issue in terms of its explanatory adequacy. I introduced my own principle *Epistemic Possibility*, and argued that it was represented the best candidate of those considered for a general account of the principle underwriting sceptical arguments from ignorance.

Looking ahead, my over-arching aim in this thesis, as set out in chapter 1, is to provide a framework for thinking about how to provide a happy-face solution to sceptical paradoxes can be found. In chapter 1, I suggested the way forward lay in pursuing two distinct, but related enterprises. The first was that of providing a diagnosis of sceptical arguments that explains how they work in terms of both sceptical hypotheses they involve, and the epistemic principles underwriting them. The second was that of providing

an error theory that explains our confusion with respect to sceptical arguments in terms that are not purely epistemological. I have undertaken the first of these enterprises between both chapter 2 and the present one. The result of chapter 2 was essentially that sceptical hypotheses can be best understood as possibilities in which it is consistent with subject believing some relevant proposition, that this proposition is false. The result of this chapter is essentially that sceptical arguments are best understood as underwritten by the principle that a subject knows a proposition only if its falsity is not consistent with them knowing what they do.

In the remainder of this thesis, I undertake the second enterprise. In chapter 6, I aim to show, ultimately, how the results of chapters 2 and 3 can be brought together with recently developed theories from cognitive psychology to suggest an error theory for sceptical paradoxes, and subsequently, a happy-face resolution of them in terms of *psychological* considerations. Before then, my first step will be to examine in detail recent approaches to providing happy-face resolutions to the sceptical paradoxes in terms of a diagnosis of, and error-theory for, sceptical paradoxes in terms of *semantic* considerations.

CONTEXTUALISM AND SCEPTICISM.

4.1 Introduction.

In chapter 1, I set out the problem of sceptical paradoxes, and the constraints on any fully satisfactory, or ‘happy-face’ resolution of them. I also argued that Pryor’s strategy, understood as a theoretical diagnosis, for responding to the sceptical problem was not a promising route to a happy-face resolution. Having developed my own diagnoses of sceptical arguments, in chapters 2 and 3, I here return to consider the details of another recent prominent approach to resolving sceptical paradoxes. Various forms of epistemic contextualism have been developed as accounts of the term ‘knows’ based on the notion of semantic context-dependence. A significant feature of epistemic contextualism is the apparent solution to the problem of scepticism that it provides.

In this chapter, I outline, compare, and contrast the suggested resolutions of three different proposed accounts of epistemic context-dependence. I begin, in section 4.2, by outlining the general nature of epistemic contextualism, and relating it to both the problem of sceptical paradoxes, and the conditions that should be met by a happy-face resolution of them. Next, in section 4.3, I run through the general contextualist strategy for resolving the paradox, and how this aims to meet the conditions on a satisfactory resolution. Finally, in section 4.4, I compare and assess the details of three distinct proposals for contextualist resolution of sceptical paradoxes. With respect to each version of epistemic contextualism, I examine the details of the proposed resolution in terms of the context-sensitivity of “knows”, and critically assess its prospects for representing a

happy-face solution to sceptical paradoxes.

4.2 Epistemic Contextualism and Scepticism.

There are a variety of differing accounts of how knowledge claims exhibit semantic context-dependence, many of which will be outlined below.¹ Underlying each is the core commitment to the view that the truth conditions of knowledge claims vary with conversational context.² In general it is maintained that ‘knows’ picks out a contextually supplied standard for counting as knowing.³ The contextualist typically holds that distinct token utterances of a knowledge sentence in distinct contexts may have distinct truth conditions as a consequence of ‘knows’ picking out distinct epistemic standards. These epistemic standards can vary across conversational context. This allows for the divergent truth values of distinct utterances of the same sentence attributing or denying knowledge in distinct contexts, even when made concurrently.⁴ For example, an utterance of “I know that Harold Godwinson was killed by an arrow” may express a truth when made in the context of a pub quiz team deliberation, owing to my meeting the low standards for knowledge that might be operating in that context. An utterance of the same sentence might express a falsehood, however, in the context of a scholarly conference on High Medieval European history.⁵

Furthermore, the contextualist allows for an utterance of a knowledge sentence, and an utterance of its negation, to have the same truth value in distinct contexts, even when made concurrently.⁶ My utterance of “I know that Harold Godwinson was killed by an arrow” may express a truth in the low-standards context of a pub quiz, even while my admission that “I do not know that Harold Godwinson was killed by an arrow” might express a truth in the distinct context of the academic conference. In this way, prior to any detailed account of how knowledge claims are context-dependent, the core commitment of contextualism can be understood as comprising both a negative and positive thesis.

¹ See (Blome-Tillmann, 2009), (Cohen, 1999), (DeRose, 1992), (Lewis, 1996), and (Neta, 2003a).

² See (Davis, 2007) and (Montminy, 2013).

³ Cf. (Williamson, 2005, 216).

⁴ See (Blome-Tillmann, 2008, 31-2), (Feldman, 2001, 61), (Greco, 2008, 417), (Montminy, 2009b, 342), and (Brady and Pritchard, 2005, 161).

⁵ There is an important sense in which they would not be exactly the same attribution or denial on the contextualist view, since ‘knows’ picks out distinct epistemic standards in the different contexts. Consequently, on that view, each utterance will express a different content.

⁶ See (Baumann, 2010, 83), (Blome-Tillmann, 2008, 32), and (Rysiew, 2001, 478). On the contextualist view, these utterances do not express contrary content because of the context-dependency of ‘knows’.

The negative thesis of epistemic contextualism is that sentences attributing or denying knowledge have no context-independent truth conditions, and so have no context-independent truth values.⁷ The positive thesis of epistemic contextualism is that sentences attributing or denying knowledge can have different truth-conditions and truth values, in distinct contexts of utterance. Both of these very closely related commitments are implied by the context-dependence of the term ‘knows’, and are indispensable to the proposed contextualist resolution of the sceptical paradoxes. Before detailing the contextualist approach to resolving the sceptical paradoxes, and the subtle roles that these commitments play in them, it will be instructive to briefly recap on the nature of the sceptical paradoxes themselves, and the adequacy constraints on their resolution.

Consider, then, the following three independently and, apparently, individually plausible, yet jointly inconsistent knowledge claims:⁸

(10) S knows that S has hands.

(11) S does not know that S is not a brain in a vat.

(12) If S knows that S has hands, then S knows that S is not a brain in a vat.

Here we have a set of claims representative of a sceptical paradox, and it will be instructive to comment on how this presentation relates to the ENVATTED BRAIN argument from chapter 1. In essence, (11) is just another statement of the minor premise, (7), of the ENVATTED BRAIN argument; likewise, (12) is simply another statement of ENVATTED BRAIN’s major premise, (8). Crucially, (10) represents the *negation* of the sceptical argument ENVATTED BRAIN’s conclusion, (9). Sceptical paradoxes are represented here, by means of (10)-(12), for the sake of ease in discussing the whether epistemic contextualist approaches represent happy-face solutions to sceptical paradoxes. This is because there are two central features of a sceptical paradox: the apparent truth and mutual inconsistency of claims such as (10)-(12). As we saw in chapter 1, in order to satisfactorily resolve the paradox, a contextualist must do two things: (a) account for the falsity of at least one of the three claims such as (10)-(12), and (b) account for the apparent truth of the three claims (10)-(12).⁹ Moreover, since identifying any one of (10),

⁷ Cf. (Schiffer, 1996, 318). On Schiffer’s way of putting things, since ‘knows’ picks out a contextual factor, such a sentence fails to express a complete proposition independently of a context.

⁸ See (Cohen, 1999, 62) and (Davis, 2004, 257-8). Cf. (DeRose, 1995) and (Neta, 2003a, 4), where the sceptical paradoxes are characterised in terms of our inconsistent ‘intuitions’ regarding the premises and conclusion of a sceptical argument.

⁹ See (DeRose, 1995, 5) and (Schiffer, 1996) for discussions of this point.

(11), or (12) as false may seem to engender some distinct unpalatable issues, it is a third condition on a satisfactory resolution that it (c) accounts for either how any unpalatable consequence of denying meeting (a) is avoided, or else why it is not, in fact, problematic.

In the rest of this chapter, I explore the manner in which the contextualist attempts to meet these conditions, in terms of a critical evaluation of the details of a variety of specific contextualist proposals for satisfying (a), (b), and (c).¹⁰

4.3 General Framework of a Contextualist Resolution.

We can set out the general contextualist framework for meeting the first two conditions on satisfactorily resolving the sceptical paradoxes. The second condition (b) can be satisfied by a contextualist resolution in virtue of the positive contextualist thesis. By allowing for knowledge sentences to have different truth-conditions, and consequently truth values, in distinct contexts of utterance, the contextualist can accommodate the truth of each individual claim (10)-(12). This is because, the contextualist maintains, the term ‘knows’ picks out a contextually determined strength of epistemic position that the subject of an attribution of knowledge must stand in with respect to a proposition p in order for an utterance of the form “S knows that p ” to express a truth in that context.¹¹

4.3.1 Accounting for the Independent Plausibility of the Individual Claims.

The contextualist can maintain that an utterance of the attribution (10) may express a truth in many ordinary contexts, where S is well enough placed with respect to the proposition that S has hands in order to meet the strength of epistemic position required by the context to truly count as knowing. This contextualist move is often explained by analogy to certain gradable adjectives, such as ‘tall’, or ‘old’.¹² For example, in some

¹⁰ Whether the general contextualist resolution can coherently satisfy these conditions at all will be addressed in chapter 5

¹¹ See, for example, (Blome-Tillmann, 2008, 31) and (DeRose, 1995, 29).

¹² See (Blaauw, 2005), (Brown, 2013, 5), (Cohen, 2000, 97), and (Feldman, 1999, 92-3). In chapter 5, I critically examine attempts to model the context sensitivity of the term ‘knows’ on gradable adjectives, or recognised indexicals. The discussion there focuses specifically on evaluating the semantic blindness objection to epistemic contextualist resolutions of sceptical paradoxes vis-a-vis happy-face solutions. My focus in this chapter is assessing, in broader terms, the prospects for providing a happy-face contextualist resolution.

everyday contexts of utterance such as a discussion about the ages of professional footballers, it may seem uncontroversial and correct to say “David is old” of the 37 year old. The player is advanced enough in age in order to meet the low standards of age required to truly count as old in that context. In a similar way, the contextualist’s positive thesis maintains that there is some ordinary context in which an utterance of (10) will express a truth in virtue of S meeting the undemanding standards for counting as knowing determined by that context.

The contextualist can maintain that an utterance of the knowledge denying sentence (11) may express a truth in other distinct contexts, where S is not well enough placed with respect to the proposition that *S is not a brain in a vat* in order to meet the strength of epistemic position required, in that context, to truly count as knowing. This contextualist move can too be explicated analogously in terms of certain gradable adjectives. For example, in some particular contexts of utterance, such as a discussion of the average British retirement age it may seem uncontroversial and correct to say “Ryan is not old” of the 39 year old.¹³ The subject is not advanced enough in age in order to meet the standards for truly counting as old in that context. In a similar way, contextualism’s positive thesis maintains that there is some context in which an utterance of (11) will express a truth in virtue of S not meeting the standards for knowledge determined by that context.

The contextualist can maintain that an utterance of the conditional sentence (12) will express a truth in any context of utterance. Since S has hands entails that S is not a brain in a vat, whatever the strength of S’s epistemic position with respect to the former proposition, it will be no stronger than the strength of S’s epistemic position with respect to the latter proposition.¹⁴ So, for whatever the contextually supplied strength of epistemic position picked out by ‘knows’, when S meets this standard for counting as knowing with respect to the proposition that S has hands, they will also meet this standard with respect to the proposition that S is not a brain in a vat.

Once again, this contextualist move can also be illustrated by analogy to gradable adjectives. For example, since the 39 year old Ryan is two years more advanced in age than 37 year old David, it will be uncontroversial and correct to say “If David is old, then Ryan is old” in any context of utterance. Whatever minimum age the standard for

¹³ In this way, ‘David is old’ represents an analogue of (10), and ‘Ryan is not old’ represents an analogue of (11). (10) is a positive categorical claim with respect to a subject *s* and the proposition that *S has hands*. (11) is a negative categorical claim with respect to *S* and a *distinct* proposition. Analogously, ‘David is old’ is a positive categorical claim, and ‘Ryan is not old’ is a negative categorical claim involving distinct individuals.

¹⁴ See (DeRose, 1995, 30-7).

counting as old is in the context, it cannot be the case that the younger subject meets this standard and the older subject does not.¹⁵ In a similar way, the contextualist maintains that in virtue of the proposition that S has hands entailing that S is not a brain in a vat, when S meets the standards for counting as knowing the stronger proposition they will necessarily meet those standards with respect to the weaker, regardless of how ‘stringent’ or ‘lax’ the standards for knowledge are within any given context.

4.3.2 Accounting for the Falsity of at least one Claim.

The contextualist’s positive thesis accounts for the independent plausibility and apparent truth of each of the knowledge claims represented by (10)-(12) when considered individually. Since these three claims are jointly inconsistent, however, the contextualist resolution must also account for the falsity of one of (10)-(12). Let’s look now, then, at how the second condition (b) might be met by contextualism in terms of its negative thesis.

The negative thesis of contextualism is that knowledge claims have no truth-conditions and, subsequently, no truth-values independently of a context of utterance. The contextualist can maintain then that which of the three knowledge claims constitutive of the sceptical paradox is false cannot be decided independently of a specific context. As I have noted, the contextualist maintains that (12) will express a truth in every context of utterance, whilst (10) may express a truth in some contexts, and so too may (11) true in some contexts. Since (12) always expresses a truth, it follows that (10) and (11) cannot both express truths in a single context. Consequently, when (10) is true (11) will be false, and vice versa.

Let C-LOW be a context with epistemic standards ‘lax’ enough that an utterance of “I know that I have hands” will express a truth in C-LOW. The contextualist maintains that some such context C-LOW exists, in which utterances of (10) and (12) would each express truths in virtue of S meeting the ‘lax’ standards for knowledge in C-LOW with respect to the proposition that S *has hands*. In this context, an utterance of (11) would express a falsehood since it denies that S meets the ‘lax’ standards for knowing with respect to the proposition that S *is not a brain in a vat*. This denial will be false since, as (12) expresses, S is at least as well placed with respect to the former as they are to the latter proposition.

¹⁵ DeRose has suggested that the truth of these conditionals is grounded in certain comparative facts: e.g. the comparative fact that “Wilt is at least as tall as Mugsy has the result that the conditional If Wilt is not tall, then Mugsy is not tall will be true regardless of how high or low the standards for tallness are set” (DeRose, 1995, 33).

Let C-HIGH be a context with stringent enough standards for truly counting as knowing such that an utterance of “I do not know that I am not a brain in a vat” to express a truth in C-HIGH. The contextualist maintains that some such context C-HIGH exists, in which utterances of (11) and (12) would each express truths in virtue of S not meeting the stringent standards for knowledge in C-HIGH with respect to the proposition that *S is not a brain in a vat*. In this context, an utterance of (10) would express a falsehood since it claims that S meets the stringent standards for knowing with respect to the proposition that *S has hands*. This attribution is false since, again, S is no better placed with respect to the latter, than they are to the former proposition.

In summary, with respect to the condition (a), the contextualist maintains that since (12) expresses an unequivocal truth in every context of utterance, there is no single context in which utterances of (10), (11) and (12) all express truths. Consequently, which of (10) or (11) expresses a falsehood cannot be settled independently of a context of utterance. In addressing condition (b), the contextualist maintains that due to ‘knows’ picking out a contextually varying standard for knowledge, utterances of (10) and (11) may express truths in distinct contexts of utterance.

4.3.3 Accounting for the Unpalatable Consequences.

In this section, I outline how the general framework of the contextualist resolution of the sceptical paradox addresses condition (c), i.e. accounting for either how the unpalatable consequences of denying (10)-(12) are avoided, or why they are not problematic. Firstly, the contextualist maintains that there is no context in which an utterance of (12) expresses a falsehood. Consequently, the *unpalatable consequence* of allowing for a subject to know that they have hands but not to know that they are not a handless brain in a vat is avoided entirely.¹⁶

The contextualist maintains that in the context in which an utterance of (10) expresses a falsehood, the unpalatable consequence of wholesale scepticism about knowledge of the external world is avoided, or at least is unproblematic, because it merely expresses that we fail to know these things only by very stringent standards for knowledge. This does not contradict our claims to know these things by the ‘lax’ standards of the ordinary contexts in which we may truthfully utter “I know I have hands”.¹⁷

¹⁶ For a discussion of this point, see (DeRose, 1995, pp. 27-33).

¹⁷ DeRose (1995, 32) explains: “on this account, the skeptic gets to truthfully state her conclusion only by raising the standards for knowledge... [which] doesn’t threaten the truth of our ordinary claims to know”. See also (Feldman, 2001, 61-2) and (Wright, 2005, 240-1).

The contextualist maintains that in the context in which an utterance of (11) expresses a falsehood, the unpalatable consequence of dogmatically claiming to know you are not a brain in a vat is avoided because of the following considerations. The semantic value of ‘knows’, and the truth-conditions of any knowledge claim, vary with a contextually determined standard a subject must meet in order to truly count as knowing. If this standard is set at a level, however low, in everyday contexts such that I meet it with respect to the proposition p without being able to conclusively prove that p , then it follows that an attribution in that context of the form “I know that p ” will express a truth.

There is an important sense in which, on the contextualist view, the charge of being dogmatic in your claims to know that you are not a brain in a vat, only bites in a context in which the standards for knowledge are stringent enough. The denial represented by (11) will express a truth in exactly those contexts where the attribution represented by (10) can only be falsely uttered. Since dogmatism is the consequence only of rejecting (11) but problematic only when (11) is true, the contextualist accounts for how the unpalatable consequence of the falsity of (11) is avoided.

For the most part, the general manner in which the contextualist proposes to resolve the sceptical paradox in terms of the notion of semantic context-dependence has been explicated above. It remains here to examine exactly how the semantic context-dependence of knowledge claims is accounted for. Numerous different versions of contextualism have been proposed in the last two decades, each of which represents a different way to put the flesh on the bones of the general contextualist resolution. In the next section, I will outline and critically assess three of these contextualist positions in terms of the conditions on a satisfactory resolution (a), (b) and (c).

4.4 Versions of Epistemic Contextualism and their Resolutions.

In the account of the general contextualist strategy for resolving the sceptical paradox detailed above, it has not been detailed how exactly it is claimed that knowledge sentences exhibit their particular epistemic semantic context-dependence. A number of different accounts of epistemic contextualism have been put forward and each proposes a unique explanation of how the contextualist resolution should be best understood.

This section outlines and assesses three of the main accounts of epistemological context-dependence: those of David Lewis, Keith DeRose, and Michael Blome-Tillmann.

In considering each account, I outline initially how the context-dependence of knowledge sentences is understood, and go on to show how this is supposed to provide a happy-face resolution of sceptical paradoxes, in terms of the conditions (a), (b) and (c). I argue that, whilst obvious progress is made in the intellectual development of the contextualist approach to resolving scepticism, each account is, ultimately, found wanting.

4.4.1 Lewis' Contextualism and its Resolution

In this section, I consider in detail the first of the three potentially happy-face resolution purportedly provided by a contextualism about the term 'knows'. This is the resolution proffered by David Lewis in his paper 'Elusive Knowledge'. In his paper, Lewis broadly characterises the nature of the sceptical paradox as being that: plausibly; we know various propositions about the external world, and knowledge is plausibly infallible, but our knowledge of external world propositions cannot be infallible.¹⁸ I think that this corresponds well enough with the characterisation of the sceptical paradox above, in terms of (10)-(12). Below, I first set out the details of Lewis' contextualist view and its supposed resolution of the sceptical paradoxes. I turn, subsequently, to critically assess whether Lewis' contextualist approach to scepticism represents a happy-face solution to sceptical paradoxes.

4.4.1.1 Lewis' Account of the Context-Dependence of Knowledge Claims.

In Lewis's account of the context-dependence of knowledge sentences, he begins with the notion underlying our verdict that (12) is true. This he identifies as the infallibility of knowledge. On this basis, he suggests a schema for knowledge sentences, such that:

(L1) "S knows that *p*" is true iff S's evidence eliminates every possibility in which not-*p*.¹⁹

In describing a subject's evidence as eliminating a possibility, *w*, Lewis is claiming that the subject's total perceptual experience and memory in *w* is not the same as what their total perceptual experience and memory is in actuality. According to Lewis, a subject's "perceptual experience *E* (or memory) eliminates a possibility *w* [just in case]... *w* is a possibility in which the subject is not having experience *E*" (1996, 553). This can

¹⁸ See (Lewis, 1996, 550).

¹⁹ (Lewis, 1996, 551). I say suggested rather than proposes because Lewis famously did not phrase his contextualism meta-linguistically.

be understood as the claim that S's evidence E eliminates a possibility w if and only if, were w to actually obtain, then E would not.

In terms of the attribution (10) and (L1), then, Lewis can maintain that if an utterance of this attribution is true of S, then S's evidence eliminates any possibility in which it is not the case that S has hands; for example, the possibility that S is a brain in a vat. If valid, this schema would secure the notion of infallibility for knowledge, and hence yield the intuitive truth of (12). However, the plausibility of the denial (11) appears to be a counter-example to the schema.

Lewis addresses this issue by expanding on the occurrence of the quantifier 'every' in the schema. The quantifier as it is used here should be subject to a domain restriction, just as it is in its innumerable other usages.²⁰ So, Lewis claims, this occurrence of 'every' should be taken to range over only the possibilities which are not irrelevant to the attribution of knowledge.²¹ For Lewis, the notion of irrelevance is tied to the notion of ignoring; in making claims about what people do and do not know, we ignore certain possibilities that we take to be irrelevant to the truth of our claims. The amended schema can now be stated:

(L2) "S knows that p " is true iff S's evidence is incompatible with every not- p possibility that is not ignored.

This amended schema provides a promising development for Lewis' account in terms of protecting the plausible truth of (10) in light of the plausible truth of (11). It could be suggested that in attributing to S knowledge of their having hands we are simply ignoring the possibility that they are a brain in a vat. Lewis points out, however, that there are further restrictions that need to be made on what possibilities are capable of being ignored or not. To illustrate the point, suppose I claim that I know where my car is, on the evidential basis that I remember parking it at the marina yesterday. It might be legitimate for me to ignore, in normal circumstances, the uneliminated possibility that it has since been dragged to the bottom of the sea, but not if I have just heard authoritative reports that a tsunami struck this morning.

²⁰ If, for example, I claimed that "Everyone is having a good time" at an awesome party, the domain of quantification might well be understood to be all and only those people who are at the party. Nonetheless, in claiming at the same party "Everyone is drunk", I may have covertly restricted the domain of quantification to the attendees who are not teetotalers, etc.

²¹ (Lewis, 1996, 553).

Accordingly, Lewis realises that his account should accommodate a distinction between when it is legitimate to ignore certain uneliminated possibilities and when it is not.²² Consequently, we can restate the schema:

(L3) “S knows that *p*” is true iff S’s evidence is incompatible with every not-*p* possibility that is not properly ignored.

With this schema in place for knowledge sentences, it is now possible to explicate Lewis’ account of the semantic context-dependence of knowledge claims. The possibilities that can be properly ignored, for the purposes of a given knowledge claim, will vary with conversational context as a function of a set of rules governing what possibilities are properly ignored or not. In other words, the semantic contribution of ‘knows’ is a relation between a subject and a range of possibilities determined by a proposition, conversational context, and a set of rules for proper ignoring.

To illustrate Lewis’ position: imagine that Brian has driven me to, and parked his car at, the local marina. Brian goes to his office, and I go to the pub. Four hours later, I have been reliably informed that a tsunami might have engulfed the local marina, but Brian has not. Aside from this, our evidential positions regarding the location of his car remain the same; we both remember that Brian parked his car at the marina. In the pub, I judge that an utterance of “Brian knows that the car is parked in the harbour” made by the pub landlord expresses a falsehood. Meanwhile, in his office, Brian judges that an utterance of the same sentence made by his co-worker expresses a truth.

Lewis’ contextualism explains the divergence in our verdicts concerning the truth values of utterances of the same sentence. On Lewis’ contextualist view, the truth-conditions of an utterance of a knowledge attribution are determined by the context of the attributor, not the subject of the attribution. According to Lewis’ view, neither Brian’s (nor my evidence) eliminates the possibility *w* = *Brian’s car has been washed away*. Since, if this possibility had obtained, our evidence would both still be the same as in the actual scenario, i.e. our memories. So, *w* remains a possibility *in which* the proposition which it is claimed Brian *knows* (that *his car is parked at the harbour*) is false.²³ Now, *in my context of assessment*, having heard what I have, it does not seem that the uneliminated possibility *w* is one that can be properly ignored. Consequently, according to (L3), I am

²² See (Lewis, 1996, 554). Lewis does this in terms of the difference between when a subject is ‘properly’ ignoring uneliminated possibilities, and when they are not properly ignoring possibilities.

²³ The possibility described by *w* (where Brian’s car has been washed away by a tsunami) is one in which the proposition we attribute knowledge of (i.e. that *the car is in parked in the harbour*) is false. In the possibility that actuality obtains, *α*, the proposition is true; his car is, in fact, parked in the harbour.

correct to judge that an utterance of “Brian knows that his car is parked at the harbour” expresses something false. After all, in the context in which I assess this attribution, Brian’s evidence does not eliminate every possibility that is *both* one in which his car is not parked in the harbour, and *not* properly ignored. In contrast, the context in which Brian assesses an utterance of the same sentence, it seems as though the possibility *w* is properly ignored. According to (L3), he may well satisfy the truth-conditions for the assessed knowledge attribution.²⁴

For the purpose of assessing the relative success of Lewis’ contextualist account at satisfactorily resolving the sceptical paradox, let us concentrate on just one of his stipulated rules governing what can be properly ignored in a context: the *rule of attention*. This rule states that in any conversational context, any possibility, *w*, that is attended to by conversational participants of the context, is not properly ignored.²⁵ By being attended to, *w* is not ignored at all. Consequently, if *w* is a possibility in which not-*p*, it becomes one of the contextually determined possibilities that S’s evidence must eliminate in order for an utterance in that context of the form “S knows that *p*” to be true.

The rule of attention is utilised by Lewis’ account to provide the details of the contextualist resolution of the sceptical paradox. By attending, in a context, to the sceptical possibility that S could be a handless brain in a vat, this possibility becomes a contextually relevant possibility. Since S’s being a handless brain in vat is a possibility in which it is not the case that S has hands, when the former possibility is attended to within a context, S’s evidence must eliminate this possibility in order for an utterance of “S knows that S has hands”, made in that context, to express a truth. However, as has been noted, S’s total evidence for the proposition that they have hands would be the same regardless of whether S has hands, or is a brain in a vat. So, when attended to, this possibility remains a salient and uneliminated possibility, in which it is not the case that S has hands. Consequently, in any context in which this sceptical possibility is attended to, an utterance of the form “S knows that S has hands” will express something false, and an utterance of the form “S does not know that S is not a brain in a vat” will express

²⁴ It could be suggested that the question of whether Brian can, or will properly ignore *w* on Lewis’ view depends on consideration of the latter’s rules of ignoring. I assume it is obvious that Brian will ignore this possibility; implicit in the story is that Brian does not believe nor entertains that *w* does or might obtain. The rules of belief, attention, and actuality do not prevent Brian from properly ignoring *w*. The rule of resemblance is interestingly different; the only obvious way that *w* saliently resembles *a* is that Brian’s evidence remains the same. Just as this makes *w* saliently resemble *a* for us, the fact that the target proposition of the knowledge attribution is false *in w*, makes *w* saliently dissimilar to *a* for us. I take Lewis to grant that the rule of resemblance does not prevent Brian from properly ignoring *w* in this case; to do so, he claims “would be capitulation to scepticism” (Lewis, 1996, 556).

²⁵ (Lewis, 1996, 559).

something true.

4.4.1.2 Lewis' Account and the Conditions on a Satisfactory Resolution.

Lewis' contextualism seems to meet the first condition on a satisfactory resolution of the sceptical paradox (a). It allows for the plausible truth of the claims (10)-(12) when considered individually. In principle, Lewis' account yields the plausible truth of (12) in any context; according to (L3), in no context can *both* an utterance of the form "S doesn't know that S is not a brain in a vat" express something false, *and yet* an utterance of the form "S knows that S has hands" express something true.²⁶ The plausible truth of the denial (11) is accounted for by Lewis' contextualism, since any claim that S knows that it is not the case that they are a brain in a vat will be false when assessed in any context in which the possibility of S being a brain in a vat is attended to. The plausible truth of the attribution (10), however, can also be accounted for by Lewis' account. Although we fail to know various external world propositions in contexts where the sceptical possibility is being attended to, the rest of the time we do not attend to these uneliminable possibilities of error, and, subsequently, we *do* properly ignore them. Consequently, in these ordinary contexts utterances of attributions such as (10) will express true propositions.

Lewis' account also appears to meet the second condition (b). It accounts for the falsity of one of the three claims, albeit in a given context of utterance. Since an utterance of (12) expresses a truth in any context, only one of (10) and (11) need be false. (10) can only be falsely uttered in those contexts in which (11) expresses a truth because the attended to possibility that S is a brain in a vat implies that S does not have hands. Since an utterance of (10) expresses a truth in those contexts where this sceptical possibility is properly ignored, an utterance of (11) will presumably express a falsehood. However, Lewis implies, according to the rule of attention, that there is no context in which it can be truly claimed that someone knows a sceptical possibility not to obtain: such a context is immediately one in which that possibility is not properly ignored.²⁷

Lewis' account seems unable to predict that an utterance of (11) can ever express a truth. Nonetheless, it might be argued that this alone does not prevent Lewis' account

²⁶ In principle, the truth of the attribution ensures there is no unignored possibility that S is a brain in a vat uneliminated by S's evidence. So, there is no impediment to S knowing that they are not a brain in a vat. There is an obvious problem for Lewis' account here. When (12) is asserted for some subject S, (10) can only be false.

²⁷ With respect to the possibility of being a brain in a vat, Lewis says "if in this context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative." (Lewis, 1996, 559). Mentioning the possibility appears to be sufficient to attend to it, and subsequently not ignore it, let alone properly.

from meeting the condition (b). After all, this condition rightly states that in order to satisfactorily resolve the paradox the falsity of at least one of the three claims needs to be accounted for, but not more. The contextualist may point out that their positive and negative theses concerning the context-dependence of knowledge claims allow them to account for both the falsity of (10) when uttered in some contexts and its plausible truth in others. Insofar as Lewis' contextualism accounts for the independently plausible truth of each claim considered individually in distinct contexts, and predicts the falsity of (10) in some contexts, his resolution undeniably meets conditions (a) and (b).

An issue reoccurs in a slightly different form when considering whether Lewis' contextualist resolution meets the condition (c), i.e. accounts for why the unpalatable consequences of any of the claims being false are avoided. Here the contextualist's reliance on their account of the falsity only of (10) in some sceptical contexts generates another issue. In order to account for why the undesirable consequence of this (i.e. scepticism about our knowledge of the external world) is avoided, the contextualist will point out that what a true utterance of the form "S doesn't know S has hands" expresses in the sceptical context doesn't contradict with what is expressed by a true utterance of "S knows S has hands" in another context.²⁸

The consequence of scepticism is avoided on Lewis' account then because when (10) expresses a falsehood in sceptical contexts, it does so because we fail to meet the unusually stringent standards for truly counting knowing determined by that context. At the same time, we do meet the 'lax' standards determined by virtually every other context. On Lewis' view, however, an utterance of (11) expresses a truth in every context in which it is considered. On Lewis' view, an utterance of (10) will express a falsehood in all those contexts in which (11) expresses a truth. So, according to Lewis, an utterance of (10) will express a falsehood in every context in which the three claims (10), (11) and (12) are considered. Lewis' contextualist is unable to coherently claim *when faced with the sceptical paradox* that an utterance of (10) expresses a truth.²⁹ Lewis is committed in

²⁸ Cf. (DeRose, 1995, 37), (Lewis, 1996, 561), and (Blome-Tillmann, 2009, 246). See also (Feldman, 1999, 96-7).

²⁹ It might be objected that this sells the contextualist resolution short. After all, the contextualist has already pointed out that we do meet the 'lax' standards in ordinary contexts when (11) isn't mentioned at all, which is not contradicted by sceptical claims made in other contexts! Lewis was, after all, not framing his account in the proper meta-linguistic manner, which lamentably gives rise to misunderstandings such as the one suggested above. I argue, however, that the problem arises exactly from framing his account in the proper meta-linguistic manner: moreover, it is only in virtue of illegitimately presenting the contextualist resolution in the object level of knowledge that Lewis' account seems to work. It is well appreciated that when Lewis speaks of us having and then losing knowledge, it is strictly erroneous but thought that the point is nonetheless grasped; we do know by low standards even if we do not by high standards. But it is not within the contextualist's remit to discuss what subjects do and do not know, over

his answer to (c) to claiming that if, in a single context, an utterance of (10) expresses a truth, then an utterance of (11) expresses a falsehood: his answer to (a) requires the situation described by the antecedent can obtain for at least some context, but his answer to (b) ensures that there is no context in which the situation described by the consequent obtains.

Lewis' account has only mixed success with its results in a couple of further explanatory areas. Since it is highly plausible that we do, in fact, know many things, a satisfactory resolution of the sceptical paradox should (d) account for why sceptical denials that involve bizarre and hard to eliminate possibilities such as the brain in a vat scenarios are more troubling than those simple sceptic denials that merely suggest that we do not know what we claim to.³⁰ It has been pointed out that Lewis' account fails here since it predicts that a claim of the form "S doesn't know that not-*p*" will ensure we attend to, and subsequently are not properly ignoring the possibility that not-*p*, which, as we saw in chapter 2 *qua* effective sceptical scenarios, it seems our evidence cannot eliminate.³¹

Furthermore, it is a further condition on any resolution that it should (e) account for our judgement that utterances of the forms "S knows that *p*" and "S does not know that *p*" in a single conversational context are contradictory.³² This is met since the semantic contribution of 'knows' to the sentence "S knows that *p*" is the same for every occurrence in a single context. Consequently, S's evidence eliminates every not-*p* possibility that is not properly ignored. Nonetheless, there remain issues surrounding the ability of Lewis' contextualism to meet each of the conditions on a satisfactory resolution of sceptical paradoxes. In particular, I think that there are reasons to doubt whether this view is able to adequately meet (c), and, consequently, constitute a happy-face solution to sceptical paradoxes.

and above discussing what an utterance of the form "S knows that *p*" will express in a given context.

³⁰ See (Blome-Tillmann, 2009, 270), (DeRose, 1995, 9), and (Neta, 2003a, 14, 21).

³¹ See (DeRose, 1995, 10-12). For discussion of both Lewis and DeRose, see (Neta, 2003a, 13-14). I am not convinced that this is a significant problem for Lewis' account, let alone the most critical. After all, there seems nothing to prevent Lewis from suggest that what prevents these 'simple' sceptical challenges from being effective is the fact that they do not satisfy the widely accepted general condition on an effective sceptical hypothesis, that they explain why you would falsely believe what you claim to know if they were to obtain. For statements of this condition see (Beebe, 2010), (DeRose, 1995), and (Neta, 2003b).

³² See (DeRose, 2004a, 3).

4.4.2 DeRose's Contextualism and Resolution.

In this section, I set out and critically assess, the second potentially happy-face resolution provided by a contextualism. This is the one put forward by Keith DeRose in his paper 'Solving the Skeptical Problem'. Just like Lewis, DeRose attempts to provide a resolution with an account of the context-dependence of knowledge sentences. This account, however, develops the contextualist strategy for dealing with scepticism by giving a central role in its resolution to the context-variant 'strength of epistemic position' a subject must have, with respect to a proposition p in order for an utterance of "S knows that p " to express a truth.³³

4.4.2.1 DeRose's Account of the Context-Dependence of Knowledge Claims.

The explanation for this contextual variation in the strength of epistemic position required to truly count as knowing, utilises a form of subjunctive conditional concerning beliefs. We have already encountered, in chapter 2, these doxastic considerations that involved in DeRose's contextualist position, and its supposed resolution of sceptical paradoxes: *Sensitivity*. We are apparently inclined to judge that a subject S does not know that some p if the following conditional is true:

(*Sensitivity*). If p were not the case, then S would not believe that p .³⁴

To have a sensitive belief that p , the subject's belief must match the fact of whether or not- p in all the possible worlds as far out as the modally closest world in which it is not the case that p . The further out in modal space one's belief matches or tracks the facts with respect to p , the stronger one's epistemic position with respect p . To have a sensitive belief that p , is to be in a relatively strong epistemic position with respect to p .³⁵

According to DeRose, in a given context, the strength of epistemic position (how far in modal space the subject's belief tracks the truth) required of a subject in order to be truly attributed knowledge varies with conversational context in virtue of the rule of sensitivity:

Rule of Sensitivity: when it is asserted that 'S knows that p ', the strength of epistemic position S must be in with respect to p for this attribution to be

³³ (DeRose, 1995, 29).

³⁴ See 2.2 and (DeRose, 1995, 18).

³⁵ (DeRose, 1995, 34).

true is raised, if necessary, to require S's belief that p to be sensitive. (DeRose, 1995, 36).³⁶

In other words, the extent to which S's belief that p must correspond to the fact of whether or not p in modal space, in order for S to truly count as knowing p can vary with context as a function of what it is asserted that S does or does not know.³⁷

In terms of the sceptical paradox, DeRose's contextualist account explains the plausibility of the sceptical denial (11) by pointing out that when uttered in a context the strength of epistemic position that S must be in with respect to the proposition that they are not a brain in a vat to truly count as knowing this proposition, can be raised to require their belief that they are not a brain in a vat to be sensitive. Their belief must correspond with the facts in the sphere of possible worlds centered on actuality that comprises all possible worlds out to the nearest world in which they are a brain in a vat. Since, as we have seen, S will still believe that they are not a brain in a vat even in a world where they are one, S's belief fails to be sensitive and consequently, an utterance of the form "S knows not a brain in a vat" is false by the rule of sensitivity, and an utterance of (11) will express a truth.

This account also accommodates the plausible truth of (10) and (12). To illustrate, consider a context in which no one is attributing knowledge of modally far-fetched propositions, but rather just claiming that "S knows that S has hands". In this case S's belief in this proposition need only match the facts in all the nearby worlds in which it is false, i.e. where they do not have hands. Since this will presumably be a world in which S either lost, or was born without hands, S will not continue to believe this in these worlds, and so, consequently, S's belief is sensitive. In this way, the rule of sensitivity yields the result that an utterance of (10) is true in this context. (12) will similarly express a truth in all contexts of utterance; if S meets the contextually determined strength of epistemic position required to count as knowing they have hands, then they will also count as knowing that they are not a brain in a vat. In all the closest worlds in which they do not have hands, S's belief that they are not a brain in a vat matches the facts.

4.4.2.2 DeRose's Account and the Conditions on a Resolution .

I now turn to discuss how DeRose's attempt to resolve the sceptical paradoxes fares with respect to meeting conditions (a), (b) and (c). To begin with, it is worth noting that it does

³⁶ Cf. 2.2.

³⁷ (DeRose, 1995, 37).

so in virtually the same generally contextualist way as Lewis' did. It does so in terms of the rule of sensitivity rather than the rules for proper ignoring. Condition (a) is met in the manner outlined above, whilst (c) is met in virtue of the proposed fact that in a context in which (11) expresses a truth and (10) a falsehood, the "very high standards [for knowledge] that we don't live up to has no tendency to show that we don't satisfy the more re'lax'ed standards [for knowledge] that are in place in ore ordinary conversations" (DeRose, 1995, 38). So, when (10) is falsely uttered, scepticism about our ordinary claims to know can be avoided.

It is not clear, however, that the consequence of dogmatism is avoided where an utterance of (11) is false. In this context the truth expressed by an utterance of (10) merely says that S's belief that they have hands tracks the truth far enough to be sensitive, and so can be truly counted as being known. The problem here is that it is in every sense question begging against the sceptic to claim that S's belief that they have hands matches the facts at all, as this assumes that actuality is a situation in which they are not a brain in a vat.³⁸ I suggest this criticism, although valid, is of little specific interest since it is central assumption of the general contextualist thesis that many of our knowledge claims can be true.

Concerning (b), DeRose's contextualism faces the same problem that confronted Lewis', which is that the difficulty in accounting for the falsity of (11). DeRose's account ties the truth expressed by an utterance of (11) with the idea that mentioning the sceptical possibility of being a brain in a vat raises the standards for counting as knowing by means of the rule of sensitivity. It is unclear whether (11) could ever be falsely uttered in any context, and, consequently, the coherence of the accounts attempts to meet (a), (b) and (c) is called into question in the manner described above.

In order to avoid this problem, DeRose claims that his account is not committed to the idea that mentioning of the sceptical brain in a vat scenario is sufficient to determine a more stringent strength of epistemic position required to count as knowing in that context.³⁹ The contextualist can maintain that the plausibility of (11) is due to the fact that it would express a falsehood in a context in which the rule of sensitivity has been utilised in the manner described above by the sceptic to "install unusually demanding

³⁸ See (DeRose, 1995, 49-50); here it is conceded that this does indeed beg the question, but that this is irrelevant to the significance of the proposed resolution. Cf. (Craig, 1989).

³⁹ This is evident throughout DeRose (1995), where we find a host of hedging terms in his discussion of how the rule of sensitivity plays the role of the mechanism by which the standards are raised (i.e. the most central aspect of his resolution): "tend" (pp. 18, 21, 23), "tend... if need be" (p. 36), "if necessary" "needn't be exception-less" (p. 19), "general inclination" (p. 40), "general tendency" (p. 40), "stopped short of fully endorsing it" (p. 44).

epistemic standards... [but w]hether the skeptic actually succeeds in [doing so]... may be a very unclear matter” (DeRose, 2006, 333). The idea here is that an utterance of (11) *could*—but need not—trigger the standards raising mechanism, the rule of sensitivity, resulting in both raised standards and truth of the denial.

In pursuing this line, DeRose’s contextualism owes an account of what situations the mere mentioning of a sceptical possibility *sh*, in an utterance of the form “S doesn’t know that not-*sh*” does result in raised standards for counting as knowing, and when it does not. DeRose attempts, in a way, to address this issue in terms of Lewis’ notion of a conversational score.⁴⁰ For DeRose, the issue depends on considerations of a single situation in which both a sceptical denial such as (11), *and* an ordinary attribution such as (10) are made by interlocutors. I argue that his account of these situations goes nowhere towards rescuing his resolution of the sceptical paradox.

Insofar as both claims are vehemently stuck to by the respective claimants, the contextualist is confronted by a puzzle. This is: how to respect the sense in which both claimants clearly indicate that their respective use of ‘knows’ is governed by divergent standards, whilst *also* respecting the sense in which they indicate themselves to be contradicting each-other.⁴¹ For current purposes, it will suffice to note that DeRose’s proposal is what he calls the gap view of such situations: where S meets both sets of divergent indicated standards, an attribution expresses a truth and a denial a falsehood, where S meets neither set, the denial expresses a truth and the attribution a falsehood, and where S meets one set but not the other, both claims lack a truth-value.⁴²

I think that the suggestion that a knowledge claim can fail to have a truth value is unhelpful for DeRose’s contextualist resolution. If S meets the standards indicated by their opponent with respect to *p*, but does not meet the standards indicated by the sceptic with respect to either the proposition that S *has hands* or S *is not a brain in a vat*, then neither an utterance of “S knows that *p*” nor “S doesn’t know that S is not a brain in a vat” will express a truth or falsehood. But I take it to be a plausible conceptual truth that either S knows *p*, or S does not know that *p*. Recall the problem that beset Lewis’ solution, namely, that the mechanism for a shift in standards is automatically triggered by the mere mentioning of the sceptical denial. In attempting to rescue his contextualism, DeRose faces some distinct issues. DeRose’s way of avoiding this issue is—at best—to advocate a semantics for knowledge sentences that inadequately captures conceptual principles about knowledge. Moreover, this account still maintains that so

⁴⁰ (DeRose, 2004b).

⁴¹ (DeRose, 2004b, 5-6).

⁴² (DeRose, 2004b, 15).

long as the sceptic does not back-down and retract their denial represented by (11), then (10) will still not express a truth when uttered.

Finally, it is worth noting that DeRose's account has the advantage of avoiding a significant weakness exhibited by Lewis' account. This concerns satisfying condition (d), i.e. explaining why sceptical denials involving sceptical hypotheses are more worrisome than their simpler counter-parts. The rule of sensitivity does not determine that latter sorts of claims require a greater degree of truth-tracking from the subject's epistemic position, and so do not result in a change in the contextual feature which the truth-conditions of knowledge claims refers to.

4.4.3 Blome-Tillmann's Account and Resolution.

In this section, I consider and critically assess the details of a third potentially happy-face contextualist resolution of sceptical paradoxes. This is the one put forward by Michael Blome-Tillmann in his paper 'Knowledge and Presuppositions'. In this paper, Blome-Tillmann argues that the best account of the contextualist resolution of the sceptical paradox is offered by some basic but important amendments to Lewis' account. These amendments are proposed with the specific intention of eliminating the serious weakness of Lewis' resolution of the sceptical paradox highlighted in section 4.4.1.2.⁴³

4.4.3.1 Blome-Tillmann's Account of the Context-dependence of Knowledge Claims.

In response to this problem, Blome-Tillmann contends that merely attending to a possibility in a context is not sufficient to render it impossible for any speaker in that context to properly ignore that possibility.⁴⁴ Rather, the rule of attention is replaced by a new rule that also stipulates certain conditions under which a possibility may not be properly ignored in a context. Blome-Tillmann's substitute is the rule of presupposition:

(RP) If a possibility w is compatible with the speaker's pragmatic presuppositions in a context C , then w cannot be properly ignored C .

⁴³ See (Blome-Tillmann, 2009, 246). Specifically, it is the failure to coherently maintain that if an utterance of (10) expresses a truth, then an utterance of (11) expresses a falsehood, and that if an utterance of (11) expresses a truth, then an utterance of (10) expresses a falsehood. It hinges, then, on the account of how (11) expresses a truth.

⁴⁴ (Blome-Tillmann, 2009, 247). To ignore that possibility either properly or not. The former obviously being the stronger; if you are properly ignoring p , then you are ignoring p . If attending to p is sufficient to make it impossible to ignore p , then it is sufficient to make it impossible to properly ignore p .

The notion of pragmatic presupposition is drawn from the work of Robert Stalnaker, and explicated by Blome-Tillmann in terms of the subsequent principle:⁴⁵

(PP) *S* pragmatically presupposes *p* in *C* iff *S* is disposed to behave, in their language use, as though they believe *p* to be common ground in *C*.

Blome-Tillmann explains the phrase ‘believed to be common ground in *C*’, following Stalnaker, by maintaining that a proposition *p* is common ground in *C* if and only if all conversational participants in *C* accept that *p*, believe that all other conversational participants accept that *p*, and so on, iterating this as far as necessary.⁴⁶ It is an important feature of (PP) that, in virtue of being linked to a subject’s own personal linguistic behaviour, the notion of pragmatic presupposition is one that allows for a given subject of attribution to determine, at least in part, whether or not they are presupposing any given proposition.⁴⁷

4.4.3.2 Blome-Tillmann’s Account and the Conditions on a Resolution.

The plausibility of the ordinary attribution (10) is easily accounted for by Blome-Tillmann’s presupposition based contextualist resolution of the sceptical paradox. In ordinary contexts, conversational participants will clearly be disposed to behave in their language use and other non-verbal actions as though they believe that the denials of sceptical hypotheses are common ground, i.e. accepted by all, etc.⁴⁸ These subjects can be properly ignoring brain in a vat possibilities and, consequently, an utterance of (10) in these contexts may express a truth, whilst an utterance of (11) would express a falsehood.

Similar considerations are employed on this view to account for the plausibility of the sceptical denial (11). In some contexts, the negations of sceptical hypothesis such as those describing brain in vat scenarios might not be believed by *S* to be common ground, and, subsequently, neither presupposed nor properly ignored by *S*. In these contexts, then, an utterance of (11) will express a truth, and an utterance of (10) will express a falsehood.

The plausible truth of (12) is accounted for in the same general contextualist manner outlined before; it is maintained to express a truth in every context of utterance. After

⁴⁵ (Blome-Tillmann, 2009, 256).

⁴⁶ (Blome-Tillmann, 2009, 250, 262).

⁴⁷ (Blome-Tillmann, 2009, 254). This will be a significant point in terms of addressing some of the issues seen to have afflicted DeRose’s contextualist account above. More on this and a potential problem for Blome-Tillmann’s account below.

⁴⁸ (Blome-Tillmann, 2009, 262, 267).

all, to truly count as knowing you have hands in any context, you must presuppose that you are not a brain in a vat (unless, of course, you can eliminate it). But, if you are presupposing that you are not a brain in a vat, then you presuppose the falsity of all uneliminated counter-possibilities to this one. As such, you can properly ignore them, and thereby count as knowing that you are not a brain in a vat.

The condition (b) is apparently met by the presuppositional contextualist account straightforwardly in virtue of how it satisfies (a). In contexts where an utterance of (10) expresses a truth, an utterance of (11) expresses a falsehood, and vice versa. Condition (c) is met on the presuppositional contextualist view too. When an utterance of (10) would express a falsehood, we avoid wholesale scepticism about knowledge because we appreciate that in other contexts, in which we presuppose sceptical possibilities, we do truly count as knowing that we have hands. The problematic consequences of rejecting (12) is obviously avoided by its universal endorsement.

When considering how Blome-Tillmann's contextualism accounts for the unappealing consequences of maintaining that an utterance of (11) could express a falsehood, the view seems to run into a problem. It is not clear that suggesting that when we presuppose that we are not in bizarre uneliminable sceptical scenarios we avoid 'begging the question' against the sceptic who says we do not know that these scenarios do not obtain. Presupposing just seems to be a way of begging the question.⁴⁹ The sceptic is plausibly interpreted as pointing out that we have no good reason to presuppose the falsity of sceptical hypotheses. The presuppositional contextualist could perhaps plausibly respond to this challenge by maintaining that this diagnosis of the situation already grants too much to the sceptic's denial. On this suggestion, the sceptic's challenge in uttering a denial such as (11) might be more plausibly interpreted along the lines of (D1) as opposed to (D2).

(D1) Consider that S might in fact be a brain in a vat.

(D2) S hasn't considered that they might in fact be a brain in a vat.

In contexts where an utterance of (10) expresses a truth, and (11) a falsehood in virtue of S presupposing that they are not a brain in a vat, if the sceptical challenge is understood along the lines of (D2) rather than (D1), the apparent problematic consequence

⁴⁹ It is important to note that the objection is not that Blome-Tillmann does not specify when it is proper or legitimate to presuppose the falsity of sceptical hypotheses, and when it is not. Rather it is that presupposing in this way, qua a resolution of the sceptical paradox just seems to amount to pretending the problem does not exist.

of ‘begging the question’ is avoided. S has considered that they might be a brain in a vat, and believes it to be common ground, nonetheless, that this possibility does not obtain.

The presuppositional contextualist appears to have only mixed success when attempting to meet conditions (d) and (e). In the case of (e), the presuppositional contextualist can maintain that the sense of intra-contextual disagreement between utterances of “S knows that p ” and “S does not know p ” could be accounted for in the following way. An utterance of the latter is true if and only if there is some not- p possibility left uneliminated by S’s evidence, which S is not properly ignoring. The former is true just in case there is no such possibility.

In the case of (d), Blome-Tillmann’s view does not offer any clear account of why it is that simple sceptical denials of the form “S does not know that p ” are less troubling as challenges to the claim “S knows that p ” than the sceptical denials involving hard to eliminate scenarios such as the brain in a vat scenario. The explanation of the plausibility of the denial (11) was that it would express a truth in a context in which S was *not* presupposing (and cannot eliminate) the possibility they are a brain in a vat. It is compatible with this that the claim “S does not know p ” expresses a truth in a context in which S is not presupposing (and can not eliminate) the possibility that not- p . There is no difference in the explanation in either case.

Importantly, however, the presuppositional contextualist account does avoid a significant problem faced by Lewis, and also, perhaps, by DeRose’s contextualist account. Unlike the latter two versions of contextualism, Blome-Tillmann’s view gives an account of how an utterance of (11) can be express a truth in some context, which does not involve the utterance itself as a factor determining the claim’s truth-conditions. It is this aspect of presuppositional contextualism that generates the apparent parity of both the simple and sophisticated sceptic noted above. This also allows this account to coherently maintain that central requirement of the general contextualist resolution, that there be a context in which an utterance of (10) expresses a falsehood and an utterance of (11) expresses a truth.

However, similarly to DeRose, in his attempt to avoid this problem by mere hedging, Blome-Tillmann owes an account of the circumstances in which an utterance of (11) would express a truth. In other words, his account needs to predict when a sophisticated sceptical challenge is effective, and when it is not. In considering the effect of the mentioning of sceptical hypotheses on conversational contexts, Blome-Tillmann distinguishes between three kinds of relevant conversational participants: ‘persistent’, ‘unsteady’, and

‘indecisive’.⁵⁰

The first consists of those who remain committed to their presupposing that sceptical hypotheses do not obtain even when they attend to these possibilities. Since persistent subjects continue to presuppose the falsity of sceptical hypotheses, Blome-Tillmann maintains that these subjects will still count as knowing the ordinary propositions. When *S* is a persistent subject then an utterance of (10) will express a truth and an utterance of the sceptical denial (11) will be false.

The second consists of ‘unsteady’ subjects. These subjects are those who instantly give up presupposing the falsity of sceptical hypotheses when these are brought to their attention by their sceptical conversational partner. As a result, Blome-Tillmann claims that these subjects now fail to count as knowing the ordinary proposition in the context. This is because *they* are no longer presupposing a conversationally salient possibility of error that is not eliminated by their evidence. When *S* is an unsteady subject, then, an utterance of (10) will express a falsehood, and an utterance of the sceptical denial (11) will express a true proposition.

The indecisive subjects are unable to decide whether to continue or stop presupposing in the way required to properly ignored, and subsequently Blome-Tillmann argues, it is unclear whether they satisfy either ‘knows *p*’ or ‘doesn’t know *p*’.⁵¹ When *S* is an indecisive subject, then, utterances of (10) and (11) will neither express truths, nor falsehoods.

These considerations suggest a way for the presuppositional contextualist to give an account to satisfy (d). Sophisticated sceptical challenges involving far-fetched possibilities of error are more troubling than simple sceptical challenges, they might argue, because they are more likely than the simpler counterparts to make conversational participants unsteady. Conversely, the simple sceptical challenges of the form ‘*S* does not know that *p*’ are more likely to incline conversational participants towards being persistent.⁵² This explanation fails to be informative and is tantamount to stating that the simple sceptic is less troublesome because people are less troubled by them.

So, what is it about the sophisticated sceptic that inclines people towards unsteadiness in the way described by Blome-Tillmann’s account? If the account can not explain

⁵⁰ (Blome-Tillmann, 2009, 264).

⁵¹ (Blome-Tillmann, 2009, 257, 265-ff.).

⁵² Blome-Tillmann tells us that it “should be obvious by now, persistent subjects are so-called because they stick to their pragmatic presupposition that $\neg sh$ when confronted with sceptical arguments, that is they remain disposed to behave, in their use of language, as if they believed $\neg sh$ to be common ground.” (Blome-Tillmann, 2009, 265).

this, then it fails to satisfy condition (d).⁵³ Perhaps, though, this is not the best way to think about the effects of persistence and unreadiness in subjects. Maybe, on this view, it would be maintained that whether a sceptical challenge is effective depends just on which sort of subjects dominate the conversational context. This route is certainly open for the presuppositional contextualist to pursue but it remains unclear whether it would bear scrutiny. At best, it seems to rest on an empirical claim that sceptical challenges will be more effective in contexts which are dominated by so-called persistent participants.⁵⁴ At worst, it is another uninformative restatement of the explanandum itself, namely that some sceptical challenges are more effective than others. Before moving on to close this section by setting out a final worry regarding the ability for Blome-Tillmann's presuppositional contextualism to satisfy condition (d), it is worth noting another familiar issue that arises here.

In attempting to provide an explanation of how an utterance of (11) could express a truth, without maintaining that the act of making such an utterance would secure this truth-value, Blome-Tillmann's account maintains that in some contexts (i.e. those dominated by indecisive subjects) knowledge claims of the form "S knows that *p*" and "S does not know that *p*" will fail to have any truth-conditions. This is the same explanatory situation that DeRose found himself in whilst attempting to give an account of the same thing, i.e. a viable resolution to the sceptical paradox based on semantic context-dependence. I suggested, at the end of section 4.4.2, that the result of proposing a 'gap view' of the semantics for knowledge sentences represents a failure to adequately capture our concept of knowledge. If this is the case, then it seems to suggest that two of the most developed contextualist resolutions of the sceptical paradox cannot satisfy conditions (a)-(e) without a tendency to fail to adequately capture our concept of knowledge. The tentative conclusion to be drawn here is, I think, that there are good reasons to think that the contextualist will be unable to provide the desired happy-face solution that we set out to find in chapter 1. According to the characterisation of happy-face solutions I outlined there, it seems that the contextualist attempts to resolve the sceptical paradoxes will, ultimately, amount to *weakly* unhappy-face solutions. This is because they appear

⁵³ A suggestion might be made in support of the presuppositional contextualist here. Recall the earlier mentioned constraint on an effective sceptical hypothesis that it explain why, if it were true, we would nonetheless continue to (then) falsely believe what we do now. If Blome-Tillmann can help himself to this, it might only take a little work to link this to the inclinations towards unreadiness rather than persistence on the part of subjects.

⁵⁴ This claim itself seems highly implausible, and I doubt Blome-Tillmann would endorse it. There is no good independent reason to deny the obvious fact that the simple sceptic's challenge are less persuasive, or effective than the sophisticated counter-parts.

to be committed, in trying to identify a false claim and give an error-theory, to an account of knowledge ascriptions that is *revisionary*.

To be disposed to behave as if one believed that it is common ground in C that *p*, amounts to a disposition to behave as if one believes that all conversational participants in C accept that *p*, and believe that all other conversational participants accept that *p*, etc. The sceptic, however, does not accept in C that the brain in a vat possibility does not obtain, even if C is otherwise dominated by persistent subjects. Consequently, it seems that no one in C can have a true belief that *all conversational participants accept that the brain in a vat possibility does not obtain*. And, so, one could only falsely believe that *it is common ground* in C that the brain in a vat possibility does not obtain.

I take it, then, that to be a persistent subject in the face of a sceptical challenge, on Blome-Tillmann's view, is to be disposed to behave as though one has a false belief (e.g. that it is common ground in C that one is not a brain in a vat). Blome-Tillmann does not give a detailed account of what sorts of things being 'disposed to behave as though you believe *p*' involves but it is reasonable to suggest that having the disposition to say "I believe that *p*" is a plausible candidate. Blome-Tillmann's account allows for a persistent subject to say "I believe that *p*" in a situation in which it is not the case that *p*, i.e. "I believe it is common ground that *q*" when it is not the case that it is common ground that *q*. The immediate issue, here, concerns the plausibility of the error theory this seems to be proposing. If I am aware that it is *not* common ground that the sceptical hypothesis is false, then, presumably I do not believe that it is common ground. As such, I will also be fully aware of that an utterance, by me, of "I believe it is common ground that the sceptical hypothesis is false" expresses a falsehood. Of course, if I am unaware that it is not common ground, then I can make my claim innocently enough.

In the former case of an aware and persistent subject, Blome-Tillmann's account predicts that I will truly count as knowing I have hands just in case I am able to falsely claim to believe something about my conversational context. In the latter case of an unaware/persistent subject, Blome-Tillmann's account predicts I will truly count as knowing I have hands just in case I am ignorant of something about my conversational context. Neither alternative is particularly appealing. The explanation of when (11) expresses a falsehood will then amount to either a conversational participant's ignorance of their context, or else their ability to utter bare-faced falsehoods about what they believe. Perhaps it could be maintained that the presuppositional contextualist account is not committed to persistent subjects being disposed to utter falsehoods such as "I believe that it is common ground that the sceptical hypothesis is false". However, if a

subject is not able to say “I believe that p ” in these situations, then it is difficult to grasp in what sense they are disposed to act as though they believe that p . In which case, on Blome-Tillmann’s view, they would no longer be presupposing that p ; rather they become unsteady, or at best indecisive subjects. In either case, on this account, (10) could not be truly uttered. And, consequently, I think there are good reasons to doubt whether either of the contextualist resolutions discussed here, are capable of providing happy-face resolutions of sceptical paradoxes.

4.5 Concluding remarks.

In this chapter, I considered another potential route by which to provide a happy-face solution to sceptical paradoxes that outlined in chapter 1. I began by outlining the general approach to resolving the sceptical paradoxes in terms of the semantic context-dependence of the term ‘knows’, and discussed the manner in which the contextualist might attempt to meet the constraints on a happy-face solution to the problem. After setting out the general contextualist approach, I next detailed three distinct accounts of the context-dependence of knowledge claims, and their respectively suggested resolutions of the sceptical paradox in detail.

I assessed the potential for each contextualist position to provide a happy-face resolution of sceptical paradoxes. I argued that the first of these accounts, proposed by Lewis, was not obviously capable of providing a happy-face resolution. I argued that two other accounts, offered by DeRose and Blome-Tillmann, also face significant obstacles to providing happy-face resolutions. My critiques of these contextualist positions drew on considerations of how the details of these different accounts of the context-sensitivity of ‘knows’ impacts on how the constraints on a happy-face solution are met. I think that the considerations raised, here, are suggestive of the inadequacy of these specific views in providing happy-face resolutions of sceptical paradoxes.

In the next chapter, I move on to discuss a broader issue with epistemic contextualist approaches to sceptical paradoxes: the semantic blindness objection.

CONTEXTUALIST ERROR-THEORIES

5.1 Introduction.

Epistemic contextualism about knowledge claims has been argued to be a promising theory in virtue of its apparent potential to resolve a number of epistemological problems and, none more so than, the sceptical paradoxes. In this chapter, I focus on discussing a single objection to the adequacy of contextualist approaches to resolving sceptical paradoxes.

In chapters 1 and 4, I highlighted that in order to satisfactorily resolve the sceptical paradox, the contextualist needs to give an account of why it is that sceptical arguments generate paradoxes.¹ An fully satisfactory, or happy-face, resolution of the sceptical paradoxes needs to explain how we got confused about by sceptical arguments. In order to be happy-face, then, a putative resolution of the sceptical paradox must *in some sense* explain why we felt a pull towards a set of mutually incompatible knowledge claims. As I noted in section 4.2, according to the contextualist semantics either the first premise of a sceptical argument, or else its conclusion will express a false proposition, depending on the context in which it is assessed. As a direct consequence, there is no “set of mutually inconsistent propositions each of which enjoys some plausibility when considered on its own”.²

If there is no genuine inconsistency, however, the contextualist is required in their

¹ See (DeRose, 1995, 3), (Schiffer, 1996, 325).

² (Schiffer, 1996, 324), emphasis added.

error theory to address the fact that we nonetheless thought that there was one. After all, according to contextualist, many of our patterns of knowledge ascriptions serve as evidence for our appreciation of the context-sensitivity of the term ‘knows’. So, it seems, that on this view we should recognise that there is no genuine inconsistency between the plausible truth of sceptical conclusions and our *ordinarily* made claims to know various empirical propositions; but yet, we clearly think there is one. The contextualist’s answer to this is that in presenting the sceptical argument, its proponent can effect or induce a rise in the standards for counting as knowing to such a level at which we no longer count as knowing practically anything. Nonetheless, the contextualist can say, we realise that the sentence expressing the conclusion of the sceptical argument *would* express a false proposition in an ordinary low standards context. So, instead we assert its negation. We do so precisely because in a low standards context the negation of the sceptic’s conclusion expresses the true proposition that S knows by low standards that S has hands.³ In this chapter, I focus on a single objection to this approach to resolving sceptical paradoxes.

I proceed as follows. In section 5.2, I introduce the semantic blindness objection to epistemic contextualist resolutions of sceptical paradoxes, and outline two different ways in which the impact of this objection can be assessed. It can be assessed either as a knockdown objection to contextualist resolutions, or as part of a comparative disadvantage on scorecard approach. In section 5.3, I concentrate on discussing the former ‘knockdown’ assessment; here, I critically evaluate a contextualist line of response, on which semantic blindness is maintained to be an expected consequence of context-sensitivity. Next, in section 5.4, I move on to consider the ‘scorecard’ assessment of the objection, and, subsequently, a contextualist line of response to it. Ultimately, I argue, the significance of the semantic blindness objection consists in the point that, insofar as contextualists are committed to semantic blindness in their error-theories for sceptical paradoxes, they are at a disadvantage compared to their invariantist opponents.

5.2 Contextualism and Semantic Blindness.

Schiffer (2004, 168-169) has noted that the contextualist explanation of why we assert the negation of the sceptical conclusion despite assenting to their premises implies a certain error theory. We erroneously assert the negation of the sceptical conclusion because we *mistakenly* think that the truth it expresses in certain contexts contradicts

³ (Cohen, 1999, 62-65), (Cohen, 2004, 191); cf. (Schiffer, 1996, 325).

our ordinary claims to know when made in low standards contexts.⁴ According to this error theory, what we fail to recognise is the most significant feature of the contextualist's semantics. That is: the true proposition expressed by the sceptic's conclusion in a high standards context does not contradict the true proposition expressed by an utterance of the same sentence in a low standards context. The contextualist's error theory is, then, apparently committed to the view that "people uttering knowledge sentences systematically confound the propositions their utterances express with the propositions they would express by uttering those sentences in certain other contexts" (Schiffer, 1996, 325).

This error theory amounts to the claim that ordinary speakers are ignorant of the semantic context-sensitivity of the term 'knows'.⁵ It implies that ordinary speakers exhibit semantic blindness with respect to knowledge claims: "their linguistic patterns fail to reflect the reality of the context-sensitivity" of knowledge claims.⁶ The sort of failure being posited here, however, is not the comparatively uninteresting sort of error that might arise if speakers fail to be aware of certain facts about the subject of the knowledge claims. To illustrate, imagine that Joe is an otherwise competent speaker of the English language, but he ascribes the property of being a US citizen to each and every Canadian citizen he meets. Obviously, each time Joe ascribes being a U.S citizen to a Canadian, he is in error. Moreover, Joe is clearly subject to a systematic sort of error in this case; he systematically mistakes Canadians for their southern neighbours. Nonetheless, this does not imply that Joe fails to understand the meaning of the terms 'citizen of the U.S.A'. The sort of error Joe is making here is not plausibly of a semantic nature. Rather, the sort of error Joe makes here is more plausibly explained in terms of his ignorance of certain relevant facts. Specifically, that some of the people to which he attributes the status of being citizens of the USA, are in fact citizens of Canada.

In contrast, however, the sort of speaker error implied by the contextualist's resolution *does* seem to be semantic in nature. The error being committed, here, according to this view, does just seem to be the error of thinking that your utterance expresses one proposition when in fact it expresses another.⁷

There is an important sense in which the contextualist's resolution is committed to the view that ordinary speakers are ignorant of the semantic context-sensitivity of the term

⁴ see also (DeRose, 1995, 5), (Neta, 2003a, 399), (Blome-Tillmann, 2008, 32).

⁵ (Montminy, 2009a, 649).

⁶ (Baker, 2012, 113). Cf. (DeRose, 2009, 159), (Hawthorne, 2004, 107), and (Schiffer, 1996, 325).

⁷ Neta (2003b, 400-4) suggests four different ways in which contextualists could explain a similar sort of error. For current purposes, I am just concerned here with spelling out what the proposed error consists in, rather than what explains/causes it.

‘knows’. Their resolution of sceptical paradoxes is committed to the view that the false claim amongst the inconsistent set cannot be determined independently of a fixed context of assessment. This is, they hold, because ‘knows’ is semantically context-sensitive. But, the objection goes, those who point to the context-sensitivity of epistemic terms such as ‘knows’ to resolve sceptical paradoxes in this way appear also to be committed to saying that most otherwise competent users of these terms are blind to this semantic fact.⁸ This is due to the sort of error theory the contextualist offers for *why* we mistakenly thought that a sceptical argument’s premises, and the negation of its conclusion, were true in a single context of assessment. Their explanation is that we got confused into thinking that each expresses a truth, by a sort of reognition that they *would* express truths in distinct contexts. But that error theory, the objection goes, amounts to maintaining that we fail to recognise—or, are blind to—the semantic context-sensitivity of ‘knows’.

5.2.1 Semantic blindness: two dimensions of assessment.

Let us consider two ways in which the impact of semantic blindness objection to the contextualist resolution from considerations of semantic blindness can be assessed.

The first is that that the objection can be understood as counting against the contextualist resolution of sceptical paradoxes on a scorecard of theoretical commitments. The idea, here, is that a semantic theory for some term *N* that does *not* posit semantic blindness concerning *N* on the part of competent users of the language in which *N* occurs, is preferable to one that *is committed* to this semantic blindness (2004, 109). On this way of assessing the objection, is that it is simply a mark against the contextualist resolution of sceptical paradoxes that it is committed to semantic blindness in its error theory. Put simply, a theory about knowledge claims should not, *ceteris paribus*, imply that people who make these claims are liable to not understand what they mean.

Attempts to defend contextualism against this objection will typically consist in arguing that whilst being committed to semantic blindness is a mark against contextualist resolutions, this does not favour any alternative invariantism about ‘knows’ over contextualism. This response concedes that the contextualist resolution is committed to semantic blindness, and that this is an undesirable feature, but argues that it is an unavoidable consequence of *any* rival semantics for knowledge claims. Some attempts to provide responses of this sort will be critically assessed in section 5.4.1, and, in section 5.4.2, will be argued to provide no help for the contextualist.

⁸ (Hawthorne, 2004, 107).

The second way of assessing the semantic blindness objection can be understood as maintaining that the consequence of semantic blindness represents a knockdown objection to contextualist resolutions. The idea, here, is that the consequence of semantic blindness itself ‘refutes’ the contextualist’s semantics for knowledge claims, and its subsequent resolution (1996, 325). On this reading, the consequence of semantic blindness is simply evidence that knowledge claims are not semantically context-dependent in the way needed for the contextualist’s resolution to work. This is because, it is argued, the consequence of semantic blindness on the part of ordinary speakers is inherently implausible for the kinds of context-sensitive terms that contextualists use to model the context-sensitivity of ‘knows’.

This objection requires further explanation since there are multiple different ways for the contextualist to model the context-sensitivity of knowledge claims. For example, in order for this objection to bear any weight, it will need to be convincingly shown that none of these options will plausibly generate semantic blindness. One possible line of response in defence of contextualism is to argue that semantic blindness is to be expected given the semantic context-sensitivity of the term ‘knows’. Again, this sort of response concedes that the semantics employed in contextualist resolution is committed to semantic blindness. However, the response denies that this is an implausible consequence of a contextualist resolution of sceptical paradoxes. I critically examine attempts to provide responses of this sort in sections 5.3.1 and 5.3.2. I argue, ultimately, that this is not a promising route for a defence of contextualism.

5.3 Semantic Blindness: knockdown objection?

In this section, I assess the semantic blindness objection as a knockdown, or refutation of the contextualist approach to resolving sceptical paradoxes. I focus on critically evaluating a contextualist response to it, on which it is denied that semantic blindness is implausible for context-sensitive terms. I consider two versions of this response, which concern two popular models for semantic context-sensitivity: indexical terms and gradable adjectives. I draw upon various diagnostic tests for both forms of context-sensitive terms, and show that ‘knows’ is not plausibly modeled on either. I conclude, consequently, that whilst this result does not demonstrate that the semantic blindness objection is a knockdown one, it is good reason to doubt whether contextualism can provide a happy-face resolution of sceptical paradoxes.

The idea that the contextualist resolution of the sceptical paradox implies that

competent speakers of the English language are systematically prone to error, explicable only by semantic blindness, has been argued to be a consequence that significantly undermines the contextualist resolution itself. The consequence of widespread and persistent error on the part of native speakers is, taken by some, to be so fundamentally implausible for context-sensitive terms that the absence of it for a given term is held to be just one aspect of a diagnostic for indexicality.⁹

However, it has also been argued that semantic blindness and the associated sort of speaker error is not implausible for other paradigmatically context-sensitive terms, such as indexical terms and gradable adjectives.¹⁰ So, taking the consequence of semantic blindness to be evidence for a lack of context-sensitivity would be a little too quick as an objection to the contextualist resolution of the sceptical paradox. Instead, I think it will be instructive to outline an argument to the effect that the theoretical consequence of semantic blindness itself renders the contextualist resolution inconsistent, and, ultimately, implausible.

The argument that the consequence of semantic blindness is a knockdown objection to contextualist resolutions of sceptical paradoxes can be understood as follows:

- (B1). It is implausible that competent speakers of the English language would exhibit semantic blindness concerning context-sensitive terms.
- (B2). The contextualist resolution of sceptical paradox implies competent English speakers exhibit semantic blindness concerning the term ‘knows’.
- (B3). The contextualist resolution of the sceptical paradox is implausible.¹¹

I think Schiffer’s objection to contextualism can be understood as implicitly advancing this sort of argument. The second premise (B2) is uncontroversial, as has been noted above, and it is worth unpacking it in more detail now. To claim, as the contextualists do, that we mistakenly assent to an utterance of a knowledge claim that expresses a falsehood in one context because it expresses a truth in another, is according to Schiffer “as though a fluent, sane and alert speaker who knows where she is, were actually to assert the proposition that it is raining in London when she mistakenly thinks she’s asserting the proposition that it is raining in Oxford” (Schiffer, 1996, 326). And this is a sort of mistake, the idea goes, which is so implausible, that it would undermine the plausibility of any theory that is committed to it as a consequence.

⁹ See, for example, (Schaffer and Szabo, 2013, 7-8).

¹⁰ See (DeRose, 2005), (DeRose, 2009), (Blome-Tillmann, 2008).

¹¹ This is my formulation of what I take to be the main argument as expressed in (Schiffer, 1996).

DeRose, however, who holds both that the term ‘knows’ is semantically context-sensitive, and that the consequence of semantic blindness is to be expected given this semantic context-sensitivity. In this way, it seems that DeRose accepts (B2), yet does not accept (B1). This is evident from his claim, amongst others, that “you can fool a lot of the speakers a lot of the time” (DeRose, 2006, 334).¹² This indicates that DeRose does not take semantic blindness concerning the context-sensitivity of ‘knows’ to be implausible, whilst Schiffer clearly does. They disagree about whether context-sensitive terms should give rise to semantic blindness. A constructive way to avoid question-begging and stonewalling on this issue, then, might be to investigate whether the context-sensitivity of the term ‘knows’ can even be plausibly modelled on other obviously context-sensitive terms.

In exploring this, I consider two forms of context-sensitivity that contextualists have used to model the context-sensitivity of ‘knows’: indexical pronouns and gradable adjectives. By examining whether the term ‘knows’ can be modelled on these terms, I think we will be in a better position to adjudicate between Schiffer’s endorsement of (B1), and DeRose’s rejection of it. To this end, in sections 5.3.1 and 5.3.2, I draw on recent work by Schiffer and Szabo (2013) to critically assess whether ‘knows’ can be plausibly modelled on either indexical terms, or gradable adjectives.

5.3.1 Denying the implausibility: the indexical model.

Concerning the alleged context-sensitivity of ‘knows’, DeRose has claimed that even when the term appears in distinct utterances of either the same, different or surface-contradictory sentences in distinct contexts, ‘knows’ “is being used with the same character, it is not being used with the same content. Or so the contextualist will claim” (1992, 921). In this vein, DeRose proposes that we might understand the context-sensitivity of ‘knows’ as an indexical term; specifically, it could be modeled on the paradigmatically indexical term ‘here’.¹³

DeRose takes there to be unavoidable semantic blindness implied by any semantic account of the sceptical paradox as well as a similar situation in which two speakers in different places respectively truly utter the sentences ‘It is snowing here’ and ‘It is not snowing here’.¹⁴ As noted above, Schiffer contends that semantic blindness is implausible in either situation. Nonetheless, I think that DeRose’s proposed modelling of ‘knows’ as

¹² A similar point is also made at (DeRose, 2009, 179).

¹³ (DeRose, 2009, 152).

¹⁴ (DeRose, 2009, 159-ff).

an indexical term could be undermined independently if ‘knows’ can be shown not to behave linguistically in the way that other indexical terms do.

Schaffer and Szabo 2013 have suggested two main tests for indexical terms: smooth tracking across contexts, and free shifting within a discourse. The former of these consists of three diagnostics for indexical terms: adjustment for indirect speech reports, the absence of cross-contextual disagreement verdicts, and the absence of widespread or systematic error.¹⁵ The third and last of the diagnostics they put forward for smooth tracking across contexts—the absence of widespread error—can be regarded as begging the relevant question, here, as to whether semantic blindness is implausible for context-sensitive terms. It is worth mentioning, however, that in the case of the indexical ‘I’, the diagnostic tells against the context-sensitivity of ‘knows’ in virtue of the indefatigable implausibility of systematic speaker error regarding the meaning of ‘I’.¹⁶

The first diagnostic for smooth tracking of indexical terms across contexts states that when an indirect speech report of an utterance involving an indexical is being given in a different context, it will require the indexical adjustment of the relevant term to produce a true report.¹⁷ For example, in order for Bill to give a true report what Sally asserted by her utterance of “I can’t wait”, his report must adjust the indexical from the first person to the second, such that it truly says ‘Sally said she can’t wait’; where a lack of adjustment yields the false report ‘Sally said I can’t wait’. The term ‘knows’ however fails this diagnostic for indexicality: no adjustment is required to yield a true indirect speech report of Jim’s utterance of “Bill knows that Sally is keen” in a distinct context, since an utterance of “Jim said that Bill knows that Sally is keen” would still be felicitous in an entirely distinct context.¹⁸

The term ‘knows’ performs badly on the second diagnostic for smooth tracking across contexts: according to which, there should be no verdict of disagreement between any two utterances involving an indexical that are made in distinct contexts. For example, if Jim utters “It is raining here” in Seattle whilst Bill utters “It is not raining here” in Madrid, there should be no judgement by anybody who is aware of where the utterances were made that Jim and Bill disagree. The contextualist resolution itself, however, maintains that inter-contextual disagreement verdicts such as this are possible. This is, after all, how they purport to explain our inclinations to both reject the conclusions of sceptical arguments, and to assent, instead, to their negations. Moreover, it is worth

¹⁵ (Schaffer and Szabo, 2013, 7-8).

¹⁶ Ibid.

¹⁷ Ibid.

¹⁸ See (2004, 101-111) for a related point regarding the Disquotational Schema for ‘Knows’.

mentioning that felicitous cross-contextual disagreement verdicts have been advocated as a diagnostic for shared and invariant semantic content in numerous places.¹⁹

The second test for indexical terms is that their semantic value can freely shift within a single discourse; in other words, if a term is an indexical, then “multiple occurrences of that... [term] in a discourse should be able to take on differing values.”²⁰ For example, it is clear that in an utterance of “Penguins live here but not here” by somebody pointing first to the south pole on a map and then to the north pole, the content of the first occurrence of “here” differs from that of the second. This is a test that ‘knows’ fails miserably; it does so for reasons related to considerations that we have already seen, in chapter 4: abominable conjunctions. Schaffer and Szabo have pointed out, for example, that if the semantic import of ‘knows’ were able to shift within a discourse, it would be possible for there to be felicitous readings of certain absurd conjunctions. One such conjunction might be, for example, the highly infelicitous sentence “Brian knows that his car is parked at the harbour, but does not know whether it has just been washed out to sea”. This sentence could be licenced, given certain contextual shifts, if the content of ‘knows’ was able to shift freely (in a way that paradigmatic indexical terms can) within a discourse.

Moreover, it is a central feature of the contextualist resolution that once the more demanding standards, by which the sceptic’s first premises is true, are in place, we could no longer truly say that “I know that I have hands”.²¹ The inability for ‘knows’ to shift in its semantic content in a single discourse is a significant point against the contextualist resolution. This is because, if this shift were possible, then the contextualist would not need to posit semantic blindness in order to explain why we reject the sceptic’s conclusion. Instead they could maintain instead that the content of ‘knows’ shifts between the premises of the sceptical argument and the speaker’s utterance “S knows that *p*” such that, in the latter, ‘knows’ picks out a relation that *does* hold of S and *p*, whilst in the former, it picked out a property that did *not*.

This move, however, not only leaves open the possibility of the absurd conjunctions mentioned above, but would also generate a verdict of disagreement of sorts: the set of claims comprising the sceptical paradox would still quite rightly strike us as inconsistent. Indeed, the availability of such a disagreement verdict can be viewed as evidence of shared, invariant semantic content. Perhaps one recourse open to the contextualist, to account for this perceived disagreement, could be to suggest that we mistakenly take

¹⁹ See, for example, (Cappelen and Hawthorne, 2009, 54-67); (DeRose, 1992, 920-1).

²⁰ (Stanley, 2005, 57); (Schaffer and Szabo, 2013, 8).

²¹ (Stanley, 2005, 66-8); (Schaffer and Szabo, 2013, 8).

there to be invariant content when in fact there is none. I find it hard, however, to see how this move would amount to more than a reaffirmation of the contextualist's error-theoretical commitment to the undesirable consequence of semantic blindness. And, ultimately, it remains unclear what help this move would be in responding to the semantic blindness objection.

In this section, I highlighted that the linguistic behaviour of the term 'knows' fails to satisfy a number of proposed tests for indexical semantic context-sensitivity. I showed that whilst paradigmatic indexical terms readily undergo the process of indexical adjustment for indirect speech reporting, knowledge claims resist this process. We saw, too, that whilst indexical terms fail to generate inter-contextual disagreement verdicts, knowledge claims infamously *do* generate these verdicts. Furthermore, I showed that whilst indexical terms are capable of freely shifting their content across multiple occurrences within a discourse, knowledge claims do not. Moreover, the last result is supported by the fact that if they were capable of this, then it would generate a strong verdict of inter-contextual disagreement, which itself would be indicative of sameness of content.

Lastly, I showed that the ability to freely shift content within a discourse would be an undesirable consequence for contextualism in virtue of licensing absurd conjunctions. I end this section, then, by concluding that there are good reasons to think that the contextualist cannot plausibly model the context-sensitivity of 'knows' on indexicals. As such, there is no route via this form of modeling to maintaining that semantic blindness is to be *expected* given the context-sensitivity of 'knows'.

5.3.2 Denying the implausibility: Gradable adjectives model.

The alleged context-sensitivity of the term 'knows' cannot be plausibly modeled on indexical terms, but another model with paradigmatic context-sensitive terms has also been advocated by some contextualists. Another model for the context-sensitive of 'knows' would be gradable adjectives such as 'tall' and 'flat'.²²

Whether the context-sensitivity of 'knows' can be adequately modeled on gradable adjective can be tested using the same methodology as was used above to argue that indexical terms do not provide an adequate model: by comparing the linguistic behaviour of 'knows' with that of paradigmatic gradable adjectives. If 'knows' were to behave linguistically in the same way as gradable adjectives, then this would support the

²² See (Cohen, 1987, 16), (Cohen, 1999, 60), (DeRose, 2005, 191), (Feldman, 1999, 92), (Hawthorne, 2004, 54), (Lewis, 1996, 554), and (Stanley, 2005, 120).

contextualist thesis, whilst a failure to do so would count heavily against any attempt to model the context-sensitivity of knowledge claims on gradable adjectives.

Schaffer and Szabo (2013) have suggested that on the first test for indexical terms—smooth tracking across contexts—‘knows’ behaves similarly to gradable adjectives. Firstly, in terms of adjustment to indirect speech reports, both ‘knows’ and the paradigmatic gradable adjective ‘tall’ resist this adjustment. Secondly, they suggest that it is at least plausible that ‘knows’ and ‘tall’ admit of cross-contextual disagreement verdicts. This is surprising in the case of purportedly gradable adjectives, since felicitous cross-contextual disagreement reports are thought to be indicative of shared context-invariant semantic content. Thirdly, however, Schaffer and Szabo do not comment on the presence of widespread systematic speaker error regarding gradable adjectives. Nonetheless, as even they recognise, these were tests for indexical terms rather than gradable adjectives.²³

The second test employed for comparing ‘knows’ with indexical terms (free shifting in content within discourse) is also applied to ‘knows’ in the comparison to gradable adjectives. On this test, it is argued that ‘knows’ also fails, since ‘knows’ does not exhibit the ability to freely shift in its semantic content between multiple occurrences within a discourse but gradable adjectives do.²⁴ Stanley uses the example of an utterance of “That butterfly is small and that elephant is small” to demonstrate that gradable adjectives—in this case, the term ‘small’—are able to appear more than once in felicitous claims despite clearly taking different semantic values at each occurrence.²⁵

As I noted above, the same feature is not exhibited by ‘knows’. Gradable adjectives also appear to be an inadequate model for the context-sensitivity of ‘knows’, since the free shifting of content within discourse is not only indicative of context-sensitivity, but rather generated by it.²⁶ Similarly to the comparison with indexical terms, Schaffer and Szabo argue that the failure of ‘knows’ to take multiple contents within a discourse precludes the ability to model the context-sensitivity of ‘knows’ that is required for the contextualist resolution to work. This is because it both prevents the content of ‘knows’ from shifting to that content by which the sceptic can truly state their premise, and from shifting, at a later stage in the discourse, to a content by which our ordinary claims to know are true.²⁷

²³ (Schaffer and Szabo, 2013, 9).

²⁴ Ibid.

²⁵ (Stanley, 2005, 57).

²⁶ See (Stanley, 2005, 58).

²⁷ (Schaffer and Szabo, 2013, 9).

I consider, now, two other novel diagnostic tests for whether a term can be plausibly modeled for context-sensitivity on gradable adjectives. The first test can be stated as being that if a term does not allow for degree modifiers, then that term is not a gradable adjective.²⁸ For example, the gradable adjectives ‘tall’ and ‘flat’ clearly allow for degree modification by the terms ‘really’ and ‘very’, as is shown by the acceptability of the sentences ‘Ben is really/very tall’, or ‘England is really/very flat’. Likewise, the negation of the degree modified adjective are able to conjoin with the unmodified phrase to produce acceptable sentences: e.g. ‘Ben is tall, but not very tall’, ‘England is flat but not really flat’.²⁹ The term ‘knows’, however, fails to do this. For example, whilst the sentence ‘Ben really knows that it is summer’ sounds acceptable, it can only be interpreted as saying that ‘Ben actually/in fact knows it is summer’, and not that he knows this to a particularly high degree of knowledge. Likewise, the sentence ‘Ben knows that it is summer but doesn’t really know it is summer’ can only plausibly be interpreted as saying that ‘Ben doesn’t know it is summer’, and not that Ben knows that it is summer, just not to a high degree.³⁰

The second test for gradable adjectives is that they should be linked to a conceptually related comparative construction: for example, ‘tall’ is linked to the construction ‘taller than’, where for any object for which there is a true predication of ‘tall’, there will natural true readings of ‘taller than’. For example, the sentence ‘The Eiffel tower is tall’ allows for the acceptable natural comparative construction ‘The Eiffel tower is taller than Ben’. However, ‘knows’ fails to allow for acceptable natural comparative constructions; the putative suggestions of ‘very much’ and ‘very well’ clearly fail to illicit acceptable sounding constructions with ‘knows’. For example, the comparative components of the sentences ‘Ben knows very well that it is summer’ and ‘Ben very much knows that it is summer’ are both more naturally interpreted non-semantically as imbuing the utterance with a pragmatic element.³¹

Furthermore, it can also be seen that ‘knows’ fails to exhibit a further third linguistic feature of gradable adjectives. This is that in the case of gradable adjectives, the comparative and degree-modified constructions are semantically interpretable with respect to a definite scale.³² For example, the sentence ‘Ben is very tall’ can be paraphrased about the extent to which figures on the height scale as in ‘Ben’s height exceeds average by a

²⁸ (Stanley, 2005, 36).

²⁹ (Stanley, 2005, 36).

³⁰ Ibid.

³¹ Ibid.

³² (Schaffer and Szabo, 2013, 10).

long way'. The term 'knows' fails to produce acceptable paraphrasing with respect to an underlying scale in this way; for example, none of the following suggestions for the completion of the sentence to include a reference to an underlying scale for knowledge seem natural or acceptable 'Ben's knowledge that it is summer is extraordinary/tough/hard to obtain'.³³

I take it that, on the picture given by these three tests for gradable adjectives, 'knows' fails to be able to be plausibly modelled on the context-sensitivity of gradable adjectives. We have seen that 'knows' again is incapable of free shift in content within discourse, and is also incapable of combining naturally with degree modifiers, or of being conceptually linked to a natural comparison construction. Finally, I have shown that 'knows', likewise, fails to allow for comparative constructions and degree modifications to be paraphrased with reference to an underlying scale.³⁴

I conclude this section, then, by suggesting that there are good reasons to think that the contextualist cannot plausibly model the context-sensitivity of 'knows' on gradable adjectives. As such, I do not think that there is a route, via this form of modeling, to maintaining that semantic blindness is to be expected for knowledge claims given the context-sensitivity of 'knows'.

So far, in this chapter, I have been considering the semantic blindness objection as a knockdown objection to contextualist resolutions of sceptical paradoxes. Ultimately, this assessment of the objection hinges on the question of whether semantic blindness is to be expected for context-sensitive terms. I presented the argument, which I take to be implicit in Schiffer, that semantic blindness is not to be expected from context-sensitive terms. I considered the contextualist response to the contrary, that 'knows' is context-sensitive, and gives rise to semantic blindness (qua sceptical paradoxes). To break the deadlock, I suggested that the onus lies on the contextualist to model the context-sensitivity of 'knows'. The bulk of this section was given over to showing that 'knows' is not plausibly modeled on either indexicals or gradable adjectives.

So, where does this leave us with the knockdown assessment of the semantic blindness objection? The implausibility of modeling the context-sensitivity of 'knows' in the above ways does not imply that the semantic blindness objection can be taken to refute contextualist resolutions of sceptical paradoxes. After all, it remains open to the contextualist to suggest that the term 'knows' exhibits an entirely novel sort of context-sensitivity. And, ultimately, I do not think that there is any plausible way in which to argue that

³³ Ibid.

³⁴ Ibid.

semantic blindness for a term is necessarily incompatible with it exhibiting semantic context-sensitivity. I conclude my discussion here by suggesting that, whilst not a knock-down objection, semantic blindness presents an obstacle to happy-face contextualist resolutions of sceptical paradoxes. Until the context-sensitivity of ‘knows’ is modeled, the contextualist *explanation* of our error with respect to sceptical arguments—necessary for a happy-face resolution—is incomplete.

5.4 Semantic blindness: scorecard approach.

In this section, I outline and evaluate the second of the two ways of assessing the impact of the semantic blindness objection against contextualist resolutions. After briefly setting out the scorecard approach to assessing semantic blindness, I move on, in section 5.4.1, to discuss what I take to be the most obvious contextualist response to this line of assessing the of sceptical paradoxes. Next, in section 5.4.2, I critically assess and reject the contextualist response to the scorecard objection. I argue that considerations of semantic blindness put the contextualist at a dialectical disadvantage with respect to sceptical paradoxes when compared with their invariantist opponents. I conclude that this suggests a promising route to finally providing a properly happy-face, fully satisfactory resolution of sceptical paradoxes.

5.4.1 Contextualist response: “you have it too”.

In defence of the contextualist resolution against the scorecard assessment of the semantic blindness objection, it could be argued that even the rival semantic accounts of ‘knows’, such as invariantism, will also have the consequence of positing semantic blindness on the part of ordinary speakers. In which case, considerations of semantic blindness alone would not count either against contextualism, or in favour of its invariantist rival.

This line of response proceeds by accepting the charge that the contextualist account of how we came to be presented with the sceptical paradox does imply that ordinary competent speakers of the English language will exhibit semantic blindness concerning the term ‘knows’.³⁵ At this point, however, the contextualist response points out that invariantism will need to posit semantic blindness in just the same explanatorily way to account for the opposite phenomena. This is because semantic blindness is posited

³⁵ See (Blome-Tillmann, 2008, 32), (Cohen, 2004, 191), and (DeRose, 2009, 159).

by contextualism to explain the verdicts of contradiction between utterances in certain cases where the contextualist thesis has committed an absence of any such contradiction.

However, the contextualist argues, semantic blindness will need to be posited by the invariantist to conversely explain the contextualist's verdicts of an *absence* of contradiction between utterances in certain cases where the invariantist thesis is committed to there being such contradiction.³⁶ The idea that a particular semantic contextualism fails to account for something in a particular discourse is not a new one. Consider a basic contextualist semantics for the predicate of subjective taste 'hilarious'. Let's imagine that on this particular imagined contextualism for 'hilarious', sentences of the form '*x* is hilarious' are construed as containing "a hidden indexical parameter which picks out the speaker's standard [of funniness]" (Baker, 2012, 109). This basic contextualist semantics for 'hilarious', suggests there is no contradiction between the certain utterances of "*Duck Soup* is hilarious" and "*Duck Soup* isn't hilarious" made in sufficiently different contexts.³⁷ This is obviously true when the utterances are made by distinct speakers. However, there is an important sense in which ordinary competent English speakers often *are* inclined to judge that these utterances are contradictory. In other words, they judge that the two speakers are genuinely disagreeing, at least, about whether *Duck Soup* is hilarious.³⁸

In order to account for the presence of these verdicts of contradiction (or disagreement) between the two non-contradictory utterances, the 'hilarious'-contextualist will posit semantic blindness. They will suggest that these agents fail to see that the term 'hilarious' is context-sensitive and the subsequent lack of contradiction. This is the same procedure that leads the contextualist about knowledge claims to posit semantic blindness in the case of the sceptical paradox. The judgement that the conclusion of the sceptical argument contradicts our ordinary claims to know, which leads us to assert its negation, is similarly explained by the contextualist in terms of semantic blindness. Moreover, the contextualist can argue that this does not favour invariantist semantics over contextualist. The idea, here, is that an invariantist will also be committed to semantic blindness to account for the fact that sceptical arguments present us with paradoxes. Along these lines, DeRose has argued that if invariantism is the correct semantics for knowledge sentences, then some sort of semantic blindness will be predicted in this case.³⁹

Following this line of thought, the contextualist will point out that *many like them*

³⁶ See (DeRose, 2009, 159).

³⁷ See (MacFarlane, 2007, 18). '*Duck Soup*' (1933) is a Marx brothers film directed by Leo McCarey.

³⁸ An interesting analysis of disagreement can be found in (MacFarlane, 2007).

³⁹ See, for example, (DeRose, 2006, 334), and (DeRose, 2009, 159).

are inclined to judge that the sceptic's conclusion that 'S does not know that p ' and our everyday claims that 'S knows that p ' do *not* contradict each other. The invariantist will maintain that either the sceptic is speaking truly when they conclude that 'S does not know that p ', or else their conversational opponent is speaking truly when they say 'S knows that p '. So, the contextualist argument goes, in the former sceptical case, the invariantist will be committed to semantic blindness to explain why we mistakenly thought that the conclusions of sceptical arguments were false. In the latter non-sceptical case, the invariantist will be committed to semantic blindness to explain why we mistakenly thought that the minor premises of sceptical arguments were false.⁴⁰

In this way, DeRose has responded to the idea that the semantic blindness objection counts as a mark against contextualism when compared with invariantism, by arguing that "there simply is a good deal of 'semantic blindness' afflicting speakers here, whichever... [semantic thesis] is correct" (2009, 160). In effect, DeRose's response to the scorecard assessment of the semantic blindness objection is to suggest to a rival invariantist that 'you have it too'. On this way of responding in support of contextualism, the argument is that semantic blindness is ubiquitous to any account of the semantics of 'knows'. Consequently, he thinks that we should simply get used to semantic blindness, rather than take it to undermine either contextualism or invariantism about knowledge claims.

5.4.2 Against the "you have it too" response.

In this section, I argue that the contextualist response to the scorecard assessment of semantic blindness is flawed. The defense, typified by DeRose's 'you have it too' line of response does not succeed in rescuing the contextualist resolution against the charge that the semantic blindness objection is a dialectical mark against contextualism compared with their theoretical opponents.

Williamson (2005) has noted that there is a certain sort of parity between epistemic invariantism and contextualism with respect to the sceptical paradoxes. The dialectic situation, with respect to the scepticism, is such that, as Williamson notes, "all theorists will be forced to postulate and explain systematic errors in our use of epistemic terms at some point or other" (Williamson, 2005, 224-225). This does not, however, imply that considerations of semantic blindness cannot be regarded as counting against either

⁴⁰ On an invariantist view, if the conclusion of the sceptical argument is false, then necessarily, one of the premises must be false given a commitment to the validity of the argument. It has already been noted that both premises are generally thought to be individually plausible.

contextualism or invariantism over the other. As Williamson explains, both theories are equally committed to positing systematic speaker error merely means that the comparative assessment of these theories should take place at the level of *how* each explains the inevitable systematic error. When faced with sceptical arguments, it appears that competent language users are subject to error, regardless of whether ‘knows’ is semantically context-sensitive, or invariant. So, which of the competing explanations for this error are the more plausible? And at this level contextualism is no longer on an equal footing with invariantism.

According to the contextualist resolution of the sceptical paradox outlined above, speakers systematically commit the error of asserting a false proposition (*S knows by high standards that they have hands*) when confronted with an utterance of a true proposition (*S doesn’t know by high standards that they have hands*). The explanation for this error is that they are unaware of the semantic context-sensitivity of ‘knows’. The invariantist, however, does not need to explain the error that speakers systematically commit according to their view by positing that speakers are unaware of the semantic context-*insensitivity* of the term ‘knows’.

On the most natural version of invariantism, which Williamson calls ‘anti-sceptical insensitive’ invariantism, the attribution of knowledge represented by the negation of the sceptic’s conclusion is true and the denial represented by the sceptic’s first premise is false. What the anti-sceptical insensitive invariantist needs to explain is the ‘illusion of ignorance’: that is, why speakers erroneously think that the latter falsehood is true.⁴¹ Importantly, the invariantist is not committed to claiming that this error is due to ignorance of the semantic workings of knowledge claims. The invariantist maintains that knowledge claims have context-invariant truth conditions such that any utterance of the sceptic’s first premise will express a false proposition in any context whatsoever.

On this view, the speaker who assents to the sceptic’s premise will not be viewed as having conflated the semantic content of the utterance with any other content it might express in another context, since this does not vary. The invariantist, then, needs to explain why the speaker mistakenly thinks that one and the same false proposition is actually true, and this does not require semantic blindness. To illustrate this point, imagine, for example, that I mistakenly believe that an utterance of “Canberra is not the capital city of Australia” expresses a true proposition. This is clearly erroneous, and we might expect some explanation to be available for my mistake. How might we explain my error in this case? It might be sufficient to explain my error to point out that I mistakenly

⁴¹ Ibid.

thought the seat of the Australian government was in Sydney rather than Canberra. This is consistent with the claim that I am aware that the sentence is true just in case a city other than Canberra is the capital city of Australia. In providing an explanation for the error of thinking a falsehood is actually true, the invariantist is free to point to non-semantic factors.

As Williamson notes, the invariantist can suggest that when faced with the sceptic's first premise and the far-fetched possibility of error that it raises, the ordinary speaker "suffers an illusion of danger" (Williamson, 2005, 226) to the effect that possibilities of error are erroneously taken to prevent us from truly claiming to know.⁴² In describing the nature of the semantic problem for contextualism about knowledge claims, Schaffer and Szabo have suggested that non-semantic factors may be taken to play a central part in generating the pattern of variable truth-value judgements which contextualists take as evidence of semantic context-sensitivity.⁴³ The point, here, is that the epistemological invariantist is *prima facie* able to explain systematic speaker error without positing semantic blindness.

An attempt to provide exactly such an invariantist account of why it is that speakers might take a literally false denial of knowledge to be true has been suggested that utilises empirical psychological data regarding belief formation and attribution.⁴⁴ Along these lines, Nagel has pointed to empirical findings that suggest that in the process of forming beliefs, the extent of cognitive effort and accuracy exhibited by subjects is directly proportional to how much is perceived to be at stake for their success.⁴⁵ A consequence of this is that "high stakes subjects pay a price from their higher accuracy, slower, and less confident belief" formation (Nagel, 2010, 283). Furthermore, empirical data reported by Nagel suggests that a subject's 'need-for closure' is broadly inversely proportional to how much is at stake for the subject to continue without this closure, (where 'closure' names the process in which a belief is fully formed).⁴⁶ Finally, the amount of evidence subjects deemed sufficient for confident belief formation was shown to be inversely proportional to the subjects' need-for-closure.⁴⁷ It could be suggested that we are less ready to attribute knowledge to subjects if we perceive them to have gone to insufficient cognitive lengths given their need-for-closure. We may find it plausible, for example, to deny that we know we are not brains in vats for the reason that given the perceived (whether reasonably

⁴² See chapter 6, where I discuss this in more detail.

⁴³ (Schaffer and Szabo, 2013, 5).

⁴⁴ See (Nagel, 2010).

⁴⁵ (Nagel, 2010, 281-282).

⁴⁶ (Nagel, 2010, 286-8).

⁴⁷ (Nagel, 2010, 288).

or not) cost of being wrong about this, we regard ourselves as having an extremely low need-for-closure, in light of which we may deem ourselves to have undergone insufficient cognitive effort to confidently form a belief that we are not brains in vats.

This is, however, only a sketch of a way in which the invariantist may account for the sort of systematic speaker error implied by an invariantist resolution of the sceptical paradox. A discussion of the empirical data used by Nagel to account for the apparent variability of knowledge ascriptions, as well as how—and whether—it should be utilised by the invariantist to account for the generation of the sceptical paradox, will be taken up in the next chapter. For the present purpose it suffices here to note that in the first instance, an invariantist account of the sceptical paradox is not committed to positing semantic blindness.

To summarise this section, I note that fact that invariantism is as equally committed to positing systematic speaker error in resolving sceptical paradoxes as contextualism does not itself place the theories on an equal explanatory footing. This is because, as I have argued, it is at least possible for the invariantist to explain this error in a *prima facie* more plausible way, through empirical data concerning belief formation and ascription, than the contextualist.

5.5 Concluding remarks.

In this chapter, I outlined the semantic blindness objection to the contextualist resolution of sceptical paradox. I considered two ways of assessing the objection. The first of these involved assessing whether the semantic blindness implied by the contextualist resolution of sceptical paradoxes amounted to a knockdown objection. I critically discussed the contextualist response that semantic blindness is not an intrinsically implausible consequence for context-sensitive terms. I drew on the work of Stanley, Schaffer and Szabo, to argue that neither gradable adjectives nor indexical terms can serve as an adequate basis on which to model the context-sensitivity of knowledge claims as required by the contextualist resolution.

The second mode of assessing the semantic blindness objection centered on the notion that, all other things being equal, it counted against contextualist resolutions of sceptical paradoxes compared with those of rival invariantists. I considered a defence of contextualism against the scorecard assessment of the semantic blindness objection, which held that semantic blindness counts equally against both invariantist and contextualist resolutions of sceptical paradoxes. I drew on the work of Williamson and Nagel, to reject

this response. Instead, it was suggested that an invariantist need not be committed to semantic blindness in their error-theories concerning sceptical paradoxes.

In the next and final chapter, I draw on the considerations above of the scorecard assessment of the semantic blindness objection, to explore the potential to provide an anti-sceptical invariantist happy-face resolution of sceptical paradoxes. I critically discuss various approaches to describing epistemologically invariant error-theories for certain patterns of knowledge ascriptions. My aim, here, will be to show how an anti-sceptical invariantist error-theory happy-face resolution of sceptical paradoxes may be found. I propose that a fully satisfactory resolution of sceptical paradoxes can be found by means of an error-theory given in terms of recent psychological accounts of our cognitive architecture.

INVARIANTIST ERROR-THEORIES.

6.1 Introduction.

This chapter develops an approach to providing a non-sceptical invariantist happy-face resolution of sceptical paradoxes. The general shape of this approach will involve using some recently developed dual process theories of cognitive psychology to provide an error theory concerning sceptical arguments.

I begin, in section 6.2, with stage setting. I outline the project to be undertaken in terms of what is involved in providing a non-sceptical invariantist error theory concerning sceptical paradoxes. Next, in section 6.3, I critically assess a recently discussed non-sceptical invariantist error theory which draws upon the resources of psychology. In section 6.4, I outline further resources of cognitive psychology that could be relevant to a non-sceptical invariantist psychological error theory. In section 6.5, I outline how the project of providing a non-sceptical invariantist error-theoretic account ‘sceptical pressure’ cases might serve as a model for a non-sceptical invariantist error theory concerning sceptical paradoxes. Lastly, I consider some potential reasons to think that the sceptical pressure error theory described might not serve as a good model for the non-sceptical invariantist to explain away sceptical paradoxes. Next, in 6.6, taking my cue from Gerken (2012; 2013), I outline recently developed default-interventionist dual-process theories of cognitive psychology. In section 6.7, I sketch a way in which the default-interventionist framework could provide a non-sceptical invariantist error theory concerning sceptical paradoxes. Finally, in 6.8, I comment on some the proposed error

theory in terms of considerations raised in previous chapters.

6.2 Semantic versus Psychological Error Theories.

Recall that sceptical arguments represent paradoxes in the sense that they seem to present us with a set of individually motivated but jointly inconsistent judgements about knowledge. The following, for example, encapsulates just such a sceptical paradox:

- (1) S does not know that S is not a brain in a vat.
- (2) If S knows that S has hands, then S knows that S is not a brain in a vat.
- (3) S does not know that S has hands.

We can distinguish between two sorts of invariantist positions with respect to sceptical paradoxes. On the one hand, an invariantist may be sceptical if they maintain that, for example, each of (1), (2), and (3) are true. In this way, a sceptical strict invariantist endorses sceptical conclusions that no one knows any empirical propositions. On the other hand, a non-sceptical strict invariantist maintains that (1) and (3) are false. As such, a non-sceptical invariantist rejects the sceptical idea that no empirical propositions are known. Sceptical paradoxes, then, present a problem for non-sceptical invariantists in the sense that sceptical denials such as (1) seem to be plausible.

A condition on a fully satisfactory, or happy-face, non-sceptical invariantist resolution of sceptical paradoxes is that it identifies the false claim involved and provides a corresponding error theory. As I noted above, on these views sceptical denials such as (1) will be identified as the false claims. Furthermore, in order to adequately resolve the paradoxes, the non-sceptical invariantist needs to explain how and why it is that we might be inclined to mistakenly think that no one knows that sceptical hypotheses (that they are a brain in a vat, for example) are true.

There are, however, constraints placed on the adequacy of non-sceptical invariantist error-theories concerning sceptical paradoxes. In chapter 1, I argued that any attempt to give an error theory just in terms of the epistemic principle underlying sceptical arguments faces higher-order problems for constituting a fully satisfactory resolution. The lesson drawn from this was that a non-sceptical invariantist needs to provide an error theory that reaches beyond *merely* identifying principles which would motivate sceptical premises. A potentially promising strategy for a non-sceptical invariantist

in this respect would be to provide a *psychological* error theory concerning sceptical arguments.

In chapters 4 and 5, I considered various epistemological contextualist strategies for resolving sceptical paradoxes. These strategies involved putting forward error theories that reached beyond epistemology in the way described. The error-theories of epistemological contextualists draw upon *semantic* considerations to explain our mistaken judgements. On these proposals, we are led into error because we are unaware of the semantic context sensitivity of the term ‘knows’ and its cognates. I argued that epistemological contextualist strategies for resolving sceptical paradoxes were not obviously happy-face resolutions. The resolutions suggested by contextualism were not satisfactory because the semantic error theories put forward failed to provide an explanation of why the sceptical paradox arose in the first place. We do not obviously make the sorts of semantic error that contextualists propose in order to account for sceptical paradoxes. And suggesting that the sceptical paradox may be considered a special case of this sort of error, leaves it an open question as to why scepticism gives rise to a paradox in this way.

The question remains whether, and how, a non-sceptical invariantist can do better than the contextualist in terms of giving an adequate error theory. In this chapter, I argue in support of the non-sceptical invariantist that an adequate error theory concerning sceptical paradoxes can be had by drawing upon the resources of cognitive psychology.

6.3 Availability and Error.

Following a suggestion of Williamson (2005), it could be argued that thinking about a sceptical hypothesis and the situation it describes, e.g. being a brain in vat, has a similar effect on us as that of exposure to fictional violence. Just as exposure to fictional violence causes viewers to over-estimate the likelihood of violent events occurring, maybe the effect of imagining scenarios in which we are deceived might be to make us over-estimate the likelihood of being so deceived. Perhaps then, the effect of considering radical sceptical hypotheses, which are vivid descriptions of global ignorance, is an over-estimation of the probability that these hypotheses are true. The suggestion here would be that we might be brought to judge that we do not know that we are not in sceptical scenarios because our knowing this would be incompatible with the probability which we now ascribe to them obtaining.¹ On this suggestion, the non-sceptical invariantist error

¹ This sort of response to sceptical puzzles was first suggested in (Vogel, 1990). For a similar suggestion see (Hawthorne, 2004). For a critical discussion of these approaches, see (Nagel, 2010).

theory concerning sceptical paradoxes would involve explaining the plausibility of the first premise of sceptical arguments in terms of the availability bias. Our assessments of the likelihood of sceptical scenarios obtaining are elevated to a level that we feel is incompatible with our knowing that they do not. The idea here is that the assessed likelihood is raised as a function of the increased ease by which we imagine these scenarios, which is a consequence of these scenarios' novelty and vividness. In this way, the non-sceptical invariantist might satisfactorily resolve the sceptical paradoxes by first maintaining that we *do* know the negations of sceptical hypotheses, and then by accounting for our tendency to judge that we do *not* in terms of the availability bias and sceptical hypotheses themselves.

I think, however, that there are some issues with foregoing approach. The suggestion that the availability heuristic—the idea that “in assessing the likelihood of an event we rely on the ease with which events of that type can be remembered or imagined” (Nagel, 2010, 290)—can help the non-sceptical invariantist explain away sceptical error has been called into question.² The ability for the availability heuristic to provide an explanation here is significantly challenged by empirical data. Recently conducted surveys, for example, show that “when a given event is difficult to imagine, subjects who have been actively encouraged to discuss or imagine it will assess it as less probable than subjects who have not” (Nagel, 2010, 292). It should be noted, as Nagel does, that the cited study found two distinct effects of the vivid imagining of a possibility on estimations of its likelihood. Where the possibility being considered was ‘easy-to-imagine’, the effect of imagining as opposed to merely reading the possibility was to increase the estimation of its likelihood. Where the possibility under consideration was ‘hard-to-imagine’ the effect of imagining as opposed to merely reading that possibility was to reduce the estimation of its likelihood.

The idea that the availability bias accounts for our tendency to judge that we do not know that we are not brains in vats, for example, involves maintaining that our estimations of the likelihood that we are brains in vats are raised by the ease with which we imagine this scenario. The empirical research suggests, however, that if sceptical scenarios involving universal deception are difficult to imagine, then considering them “in vivid and convincing detail” (Williamson, 2005, 226) would result in our *lowering* our estimations of their likelihood. Given the empirical data, the suggestion that imagining sceptical scenarios causes us to over-estimate their likelihood would imply that these

² For more on the availability heuristic, see (Tversky and Kahneman, 1973) and (Tversky and Kahneman, 1974)

scenarios are ‘easy-to-imagine’ rather than ‘hard-to-imagine’. It is not clear, however, what will serve as good evidence for the claim that sceptical scenarios—for example, that one is a brain in a vat, or currently dreaming—can be considered easy to imagine rather than hard to imagine. And in this way, it is difficult to assess the plausibility of this suggestion in these terms alone.

There are, however, some other considerations which I think bear importantly on whether the availability heuristic is a helpful way for a non-sceptical invariantist to explain our error with respect to sceptical paradoxes. It has been noted that any account of the plausibility of sceptical denials involving sceptical hypotheses should also account for the corresponding lack of persuasiveness of the so-called *simple sceptic* who “simply insists that you do not know that you have hands, offering no reasoning at all for this skeptical assertion” (DeRose, 1995, 9). On the availability account, the ease of imagining sceptical scenarios increases our estimations of their likelihood to a level that is deemed incompatible with knowing they do not obtain. But we do not want to say that the ease of imagining simple scenarios increase our estimations of their likelihood to a level that is incompatible with knowing them not to obtain.

The proponent of the availability account needs to explain the apparent fact that our estimations of the likelihood of the simple scenario’s obtaining are not raised in this way. There appear to be two ways of going about this, but neither is without issues. On the one hand, it could be maintained that ease of imagining results in a significantly *greater degree* of over-estimation of the likelihood of sceptical scenarios, compared with that of simple scenarios. On this picture, though the the estimated likelihood of the simple scenario increases as a function of ease of imagining, it does not rise to a threshold deemed incompatible with knowing it does not obtain. In the the absence of any explanation for why this might be, however, this suggestion appears *ad hoc*.

On the other hand, the proponent of the availability account could maintain that imagining simple scenarios does not result in an increase in our estimations of their likelihood at all, but rather a decrease in the way described above. This would imply that they are more difficult to imagine than the sceptical scenarios whose likelihood we do over-estimate. It would be difficult to motivate the claim that simple hypotheses are harder to imagine than the sceptical hypotheses. The claim that “exposure to lurid stories about brains in vats, evil demons, painted mules, or gamblers who bet the farm” generates an over-estimation of the likelihood of these scenarios obtaining, in virtue of the availability heuristic, appears to be committed to the implausible claim that simple descriptions of error generate an even greater sense of “epistemic danger” (Williamson,

2005, 226). In which case, an availability heuristic account of our negative epistemic judgements concerning the negations of a sceptical hypotheses does not seem to account for the lack of any corresponding negative epistemic judgements concerning the negations of simple alternatives.

Additional empirical data also presents issues for the suggestion that the availability heuristic can help a non-sceptical invariantist explain our error with respect to sceptical paradoxes. Nagel discusses studies that appear to show that “spontaneous discounting”—an “automatic tendency to discount the significance of availability whenever an offsetting explanation of increased availability” (Nagel, 2010, 295)—can be expected to cancel the effect of the availability heuristic. One cited study, for example, revealed that participants who had not been primed to consider the initial letters of words over-estimated the frequency of the first letters of their names in a given English language text; this result demonstrates the availability heuristic. In contrast, this study also found that participants who had been primed to consider initial letters, by being asked to write their initials and read a passage that concerned the significance of initial letters, under-estimated the frequency of their own name’s initial letters (Nagel, 2010, 297). This result shows that the over-estimation effect of the availability heuristic is *cancellable*. Moreover, the effect of availability is significantly lessened when the factors generating availability are made more overt to the participants.

If the availability suggestion were correct, then this data suggests that entertaining a *more* explicit, or ‘lurid’ sceptical hypothesis will result in a *greater* discounting effect. The availability suggestion holds that consideration of sceptical hypotheses, such as the brain in a vat hypothesis, induced over-estimations of epistemic danger through ease of imagining. The empirical studies cited by Nagel, however, suggest that this effect will be offset by the discounting effect. That is, “some alternative causal hypotheses for availability might be so salient as to alert people to the potential bias of using availability in judgment. Such circumstances not only might invoke discounting automatically, but also might cause participants to overcorrect for potential bias, leading to bias in the other direction” (Oppenheimer, 2004, 101). The point here is that if the novelty and strangeness of sceptical hypotheses were responsible for any availability effect, then empirical data predicts that the salience of these factors would also result in our adjusting for, or ‘discounting’, the availability effect. Consequently, the availability suggestion is at odds with what we would predict on the basis of empirical data. We would not expect, for example, to come to mistakenly judge that we do not know we are not brains in vats on the basis of this possibility being made salient.

To provide a happy-face resolution of sceptical paradoxes, the non-sceptical invariantist needs to provide a better explanation for the error we might make in judging that we do not know that sceptical hypotheses are false. Williamson's suggestion that such an explanation could be given in terms of the availability heuristic has been argued to face significant difficulties. In the rest of this chapter, I explore the possibility that a more plausible explanation of our erroneous judgement may be found by investigating the architecture of the cognitive processes that generate our epistemic judgements. In the next section, I briefly outline some of the main aspects of work within psychology concerning cognitive processes underlying our judgements about knowledge ascriptions. I then move on, in section 6.5, to survey work by Nagel, who draws on insights from cognitive psychology to provide a non-sceptical invariantist error theory concerning a related pattern of judgements about knowledge ascriptions.

6.4 'Mindreading' and Dual Process Theories.

In this section, I outline some of the resources from cognitive psychology that can be brought to bear on the sketch of a non-sceptical invariantist error theory concerning the sceptical paradox. My aim in this chapter is to draw upon these resources to put forward a non-sceptical invariantist error theory which may be offered as an explanation for why sceptical arguments can be paradoxical. I offer a brief account of the main aspects of this work which I take to be relevant to the project of providing a non-sceptical invariantist error theory. I begin by introducing the 'mindreading' capacity, then move to the 'dual process theory' of judgement, and, ultimately, the more specific 'default-intervention' model of dual process theory.

The capacity to attribute mental states to individuals has become broadly referred to as the 'mindreading' capacity in the recent literature in psychology and philosophy.³ The importance of this capacity in our everyday lives is hard to understate. As Peter Carruthers puts it, we "are inveterate mindreaders. We routinely (and for the most part unconsciously) represent the mental states to [sic] the people around us... We attribute to them perceptions, feelings, goals, intentions, knowledge, and beliefs, and we form our expectations accordingly" (Carruthers, 2009, 121).⁴

³ For comprehensive discussions of mindreading, see (Nichols and Stich, 2003) and (Goldman, 2006).

⁴ It is worth noting that theorists distinguish two levels of mindreading. The attribution of propositional attitudes such as belief and knowledge are considered to be 'high-level', whilst those of perceptions, feelings and desires are 'low-level'. See chapters 6 and 7 of (Goldman, 2006).

The relevance of the mindreading capacity to the problem of sceptical paradoxes for a non-sceptical invariantist is readily appreciable. The non-sceptical invariantist needs an account of why we might mistakenly judge that we do not know that we are not in sceptical scenarios (for example, where we are dreaming, or brains in vats). What needs to be explained is why we are systematically misled with respect to a certain class of knowledge attributions, viz. those involving sceptical hypotheses. Insofar as knowledge attributions are attributions of mental states, then it is reasonable to view these erroneous judgements as products of the mindreading capacity. In seeking to explain why we are systematically mistaken, I think a non-sceptical invariantist will do well to appeal, in presenting their error theory, to facts about our mindreading capacities. Here, I want to briefly outline some aspects of the mindreading capacity that might be relevant to the project of explaining why our judgements about knowledge go awry with respect to sceptical scenarios.

The mindreading capacity is a *cognitive* capacity, since, as Alvin Goldman explains, “[a]ttributing mental states is forming beliefs about their tokenings” (2006, 9). I take it then that an account of the cognitive basis underpinning this capacity will be directly relevant to the task of explaining how and why we might come to form an erroneous belief via this capacity (e.g. that we do not know that sceptical hypotheses are false). The simple point to be stressed here is that an explanation for the error here will involve a description of how an attributor’s cognitive processes function to produce an erroneous output given the relevant input of sceptical hypotheses.

There is widespread support for the view that the architecture of human cognition in general, and subsequently, the cognitive bases of our judgements about knowledge are captured by so-called dual process theories.⁵ All dual process theories distinguish between at least two distinct types of cognitive processes which underlie our judgements. Following Evans (2014), I will use the labels ‘Type 1’ and ‘Type 2’ processes; Type 1 processes may also be referred to as ‘intuitive’, whilst Type 2 processes can be called ‘reflective’. Type 1 and Type 2 processes are distinguished by sets of typical features. On the one hand, Type 1 processes are typically correlated with being fast, non-conscious, and automatic. In contrast, Type 2 processes are typically correlated with being slow, conscious, and deliberate.⁶

⁵ For example, (Evans, 2009), and (Tversky and Kahneman, 1996). See (Frankish, 2010) for a comprehensive overview of the development of dual process theories. For opposition to dual process theories of cognition, see (Erana, 2012). For discussion of the most prominent alternative to dual process theories of mindreading, see (Herschbach, 2015).

⁶ For a detailed discussion of these features see (Evans and Stanovich, 2013; Evans, 2014), and (Stanovich and Toplak, 2012).

Above I suggested that dual process theories could help illuminate an explanation for our mindreading capacity to err in delivering the judgement that we do not know that sceptical hypotheses are false. An important aspect of my account will be to acknowledge the sense in which internal conflicts can arise when making knowledge judgements.

“Dual-process approaches do not simply account for why we reason effectively on some problems and poorly on others; they also capture the *internal conflicts* that arise in judgement and reasoning tasks, when people’s analytic processes clash with their heuristically based intuition” (Stupple et al., 2013, 55, my emphasis).

The relevance of so-called dual-process theories of mindreading to the sceptical problem can be made more explicit. The working hypothesis of this thesis is that sceptical paradoxes are precisely such ‘*internal conflicts*’ between our reflective, philosophical judgements, and our intuitive, ‘ordinary’ judgments about what we know. Sceptical arguments are paradoxical, in the sense that they elicit seemingly inconsistent judgements from us. In order to satisfactorily resolve these paradoxes, we need to explain why we are led to make these systematic errors in our judgements by sceptical arguments. We have already analysed sceptical arguments in detail, and we know that the paradoxes are a function of these arguments, and how we think about knowledge. The only thing missing from our error theory is then a *description* of how the cognitive processes underlying our judgements about knowledge work, which sheds light on how we end up in these paradoxical situations. And this is what dual-process theories of mindreading can provide us with. Before outlining my preferred account of how a non-sceptical invariantist can provide an error theory to account of sceptical paradoxes, I move on in the next section to critically assess a potentially promising approach.

6.5 Invariantism, Error and ‘Sceptical pressure’

In this section, I describe a model for a non-sceptical invariantist error theory with respect to sceptical paradoxes. And I begin, now, by describing Nagel’s project concerning a pattern of judgements about knowledge ascriptions involving ‘sceptical pressure’ cases. Sceptical pressure cases are vignettes which typically elicit a denial that an agent knows a proposition. They differ from vignettes which typically elicit knowledge attributions to the same agent and proposition only in terms that strict invariantists’ maintain are irrelevant to the truth of the attribution. My hypothesis here is that a strict invariantist

account of these cases will be directly related to sceptical paradoxes, and I motivate the use of Nagel's account of sceptical pressure cases as a model for a non-sceptical invariantist account of sceptical paradoxes. Next, I outline the general shape of the error theory described by Nagel with respect to sceptical pressure cases. Finally, I describe a putative account of how a non-sceptical invariantist might apply the framework of Nagel's sceptical pressure case error theory to sceptical paradoxes.

6.5.1 'Sceptical Pressure' Puzzle.

Jennifer Nagel has conducted experiments in which some participants responded to—amongst other things—*one* of the following three cases that describe the same subject in relatively similar environments who is presented as making a judgement about the time:

“(A) Wanda is out for a weekend afternoon walk. She lives in a large new condominium tower down-town, and her suite is fairly small and does not have any windows that open, so she really likes to get out for some fresh air. Passing near the train station, Wanda wonders what time it is. She glances up at the clock on the train station wall and sees that it says 4:15 pm. It is in fact 4:15 pm at that moment.

(B) Wanda is out for a weekend afternoon walk near the train station and wonders what time it is. She glances up at the clock on the train station wall and sees that it says 4:15 pm. It is in fact 4:15 pm at that moment. The station clock is in fact working, but it has no second hand, and Wanda only looks at it for a moment, so she would not be able to tell if the clock were stopped.”⁷

The second of these cases (B) is a sceptical pressure variant of (A), and differs from the latter only in that it “raises the possibility that the clock is stopped, underscoring the fact that nothing in the protagonist's evidence specifically excludes that possibility” (Nagel, 2012, 174). Only 58 percent of participants responding to (B) judged that Wanda knows that *it is 4:15 pm*. Across the total eight versions tested, the subjects of the sceptical pressure cases were judged to be cases of knowledge only 39.8 percent of the time (Nagel, 2012, p.175). A much higher percentage (72.0) of participants judged that

⁷ (Nagel, 2012, 174). These are three variants of the Wanda story. In the experiment, each participant responded to a series of different stories, which included only one variant of the Wanda story, i.e. only one of (A) or (B).

the subject of the 'ordinary' cases, including (A) and its variant story counterparts, knew that it is 4:15.⁸

Given this data, it remains to be explained what it is about the sceptical pressure case that generates a negative epistemic judgement concerning the subject. How does the difference between the cases (A) and (B) generate the divergence in subjects' willingness to ascribe knowledge in the two cases? My interest will be in assessing whether the approach Nagel takes to accounting for *this* seemingly inconsistent pattern of knowledge ascriptions for an invariantist could serve as a model to explain the seemingly inconsistent pattern of knowledge ascriptions constituting sceptical paradoxes.

6.5.2 Explaining the Error.

How might a non-sceptical invariantist explain the problematic pattern of judgements about knowledge ascriptions with respect to sceptical pressure cases? Nagel maintains that an implausible explanation would be that participants judged that "a momentary look at a clock cannot be a way to know the time... [since] case (A) also stipulated that Wanda just "glances up at the clock" and yet 86 percent of participants who responded to (A) judged that Wanda did know the same proposition (Nagel, 2012, 175-6).⁹

To begin providing an alternative non-sceptical explanation of the propensity for negative epistemic assessments of subjects in sceptical pressure cases, Nagel highlights an important psychological fact. This is that, when we are attempting settle some question "we have various more and less elaborate ways of making up our minds" and these "various more or less elaborate ways of thinking about a problem are known as 'cognitive strategies' in the literature on variations in cognitive effort" (Nagel, 2012, 178). These can be broadly classed into two distinct types: "...low' strategies are quick and heuristic in character; 'high' strategies demand greater effort and deliberate sequential consideration of various alternatives" (ibid.). Nagel explains that a crucial aspect of both story variants (A) and (B) is that none of these cases describe or provide information concerning the cognitive strategy Wanda is employing in her attempt to find out what time it is.

This is a significant omission since Wanda's judgement as to what time it is will be a result of both her cognitive processes and her evidence. Yet our ability to correctly

⁸ There were no reported p-values for the study.

⁹ Of course, a sceptical invariantist could argue that this is the right explanation. They could maintain the 86 percent of participants attributing knowledge to Wanda in (A), and the 42 percent that did so in the case of (B) are mistaken.

represent the mental states of others—our mindreading capacity—is limited by the fact that “we cannot simultaneously execute two strategies, or resolve a problem in two ways at once” (ibid). The point here is that only one of at least two (broadly construed) possible cognitive strategies can be employed at any one time. So, if the cognitive strategies that we are using to think about a particular question diverge from those of an agent whose judgement we are assessing, then we will be unable to accurately represent the strategy which *they* employ in *our* assessment of their judgement.

The reason for this, Nagel explains, is that our capacity to ascribe mental states to others takes as inputs only those “outwardly detectable features of their evidential position” (Nagel, 2012, 179), and not any information concerning what cognitive strategy they are in fact employing. When taken together with the thought that “Wanda’s thinking *should* be construed the same way” in each of the two cases (A) and (B), Nagel suggests that the divergence in our epistemic assessment of her beliefs in each case is a consequence of the fact that we may not “*intuitively see* Wanda as thinking the same way when we read each [vignette]” (Nagel, 2012, 180, added emphasis). The question is, then, how do the two cases generate different ways of seeing the cognitive strategy Wanda employs in the forming her judgement about the time?

Nagel explains that a central area of agreement in the literature on mindreading includes the idea that “our mindreading system works on the assumption that given the same inputs, people will tend to think in the same way”, which is to say that “a person’s cognitive strategy is intuitively represented or replicated *just on the basis* of an appreciation of her inputs” (Nagel, 2012, 181). The point here is that since the subject’s inputs are identical in both (A) and (B), the subject will intuitively be seen to be pursuing the same cognitive strategy in each. Nagel’s suggestion is that our divergent epistemic assessments between the two cases may reflect the fact that features of the cases lead to the employment of different cognitive strategies “*as a standard* or benchmark for the performance of the observed agent” (Nagel, 2012, 183). The point here is that if a subject’s judgement is assessed against the standard of the cognitive strategy employed by the attributer, then if the subject’s epistemic behaviour is seen to fall short of that expected of a subject pursuing the attributer’s cognitive strategy, then the subject will be seen as having a problematic basis for their judgement.

How might an attributer responding to cases come to employ a different cognitive strategy from the subject of the cases? Nagel points out that an important feature typical of sceptical pressure cases is that they “invite us to contemplate merely potential inputs—evidence that could have been collected [by the subject] but was not” (ibid). In

case (B), for example, the reader of is informed that *had* the clocked been stopped, Wanda *would* not have been able to tell that this was the case. Nagel explains that an important consequence of this counter-factual information is that in order to “understand the last negated clause” the reader of the case needs to think about the “possibility in which Wanda *is able* to tell whether the clock is stopped” (Nagel, 2012, 183-4). Thinking about this possibility involves thinking about the subject doing more than they actually do, in coming to form their judgement. The subject in (B), for example, could be thought of as spending more time looking at the clock—presumably 60 seconds would be sufficient—since this is what we might expect it would take for them to be able to tell that the clock had stopped. In order to appreciate that Wanda *would* not be able to tell whether the clock had stopped, we need to appreciate the extent to which her epistemic behaviour falls short of *what it would take* for her to do so.

The most significant point here is that the *counterfactual* information involved in understanding the sceptical pressure cases draws our “attention to evidence that might have been collected but was not” (ibid). It allows us to appreciate a gap between what it is that the subject in fact does and what they *could* do further in forming an accurate belief about what time it is. Importantly, thinking about “the hypothetical collection of surplus information” by the subject of the case “requires the representation or replication of a more elaborate cognitive strategy” by the reader than actually employed by the subject (ibid). Our consideration of the possibility of error in reading the case has the effect of pushing us “into a higher cognitive strategy for the problem the subject faces: we move from making an automatic judgement to reasoning explicitly or sequentially about the reliability of our source” (Nagel, 2012, 186). Importantly, in this explanation, Nagel is alluding to the dual-process theory distinction between automatic, effortless, fast Type 1 or ‘intuitive’ judgements, and conscious, effortful, slow Type 2, or ‘reflective’ judgements. This parallels my suggestion above that an adequate non-sceptical invariantist treatment of sceptical paradoxes ought to draw explicitly on this distinction.

Nagel suggests that since we evaluate the fit between their thinking and their objective environment when we assess whether a subject has knowledge, then the standard by which *we* assess their thinking about their environment may plausibly be our *own* way of thinking about *their* environment. Since the consideration of possibilities of error induce the adoption of a more complex, effortful strategy by the attributer than that understood to be pursued by the subject, Nagel suggests an explanation for the general reluctance of attributers to ascribe knowledge to subjects in sceptical pressure cases. The suggestion is that this reluctance to attribute knowledge represents our

censuring of their “failure to adopt either our cognitive strategy or the range of evidence we now find intuitively necessary [to form a reliable judgement], given the strategy we have adopted” (ibid). The subject’s epistemic behaviour in the sceptical pressure case (B) is the same as in the ordinary cases such as (A), in which knowledge is generally attributed. An invariantist will maintain that if this epistemic behaviour is sufficient for the subject to form a reliable belief, and subsequently to know the relevant proposition in the latter case, then it is similarly sufficient in the former case.

Now, a *sceptical* invariantist could maintain, conversely, that the subject does *not* know the relevant proposition in *either* the sceptical pressure cases, *or* the ordinary cases. Yet an invariantist need not be sceptical in this way, and it is clear that Nagel does not think that our reluctance to attribute knowledge in sceptical pressure cases motivates sceptical invariantism. This reluctance, on her account, is simply the result of the attributor representing the subject as failing to employ a more cautious belief forming method than the one they really employ. But that does not mean that it would be correct to *deny* that the subject knows, since, as Nagel puts it “[u]nless we have already been persuaded by the skeptic, we do not ordinarily think that people should only ever form their beliefs cautiously and inferentially” (Nagel, 2012, 187). A sceptical invariantist could maintain that a denial that the subject knows in the sceptical pressure case is correctly responding to the failure of the subject to respond to genuinely problematic basis for their judgement. They could also maintain that an incorrect attribution of knowledge, in the ordinary case, is simply the result of failing to respond to the *same* problematic basis of the subject’s judgement. Nagel’s account of our intuitive epistemic assessments of the two cases leaves it open whether subjects of ordinary cases “fail to try as hard as we [recognise they] should” in order to know, or whether the subjects of sceptical pressure cases are erroneously expected by us to expend “extra effort [in order to know] needlessly” (Nagel, 2012, 187).

To summarise this section, it will be useful to bring together some of the most important takeaway points from the discussion of Nagel’s work on sceptical pressure. Central to Nagel’s work described here was the attempt to provide a non-sceptical invariantist error theory, or an account for a seemingly problematic pattern of epistemic judgements. Specifically, those elicited in response to *both* sceptical pressure cases, and their ‘ordinary’ variants. In line with the concluding remarks of section 6.3, my interest here has been in the fact that this error theory is given in terms of the underlying cognitive processes generating our epistemic judgements. The distinction between reflective and intuitive judgements was seen to be central to the account of the erroneous epistemic judgement

in response to sceptical pressure cases. Can the same account be used to explain the erroneous epistemic judgement in response to sceptical arguments? In what follows, I assess the extent to which the dual process error theory Nagel offers for sceptical pressure cases is applicable to the sceptical paradoxes.

6.5.3 Dual Processes, Error and Sceptical Paradoxes

In this section, I offer a putative account of how the error theory proposed by Nagel for sceptical pressure cases might apply to sceptical paradoxes. Central to Nagel's error theory was the effect of counter-factual information on the reader's cognitive processes. Specifically, the effect of considering the counter-factual information presented in the sceptical pressure case was to require the reader to engage in reflective, Type 2 processing. On Nagel's account, the erroneous judgement that the subject does not know the target proposition in response to the sceptical pressure case, in which counter-factual information is made salient, is a direct effect of this shift in how the reader thinks. To reiterate, the effect of the counter-factual information on the reader is to get them to think reflectively about epistemic behaviour bearing on the target proposition *which the subject has not pursued*. This subsequently generates the reader's judgement that the subject *does not* know the target proposition.

Perhaps a promising way for a non-sceptical invariantist to account for our erroneous judgment in the case of sceptical paradoxes would be in similar way to how they account for the sceptical pressure cases. The latter sort of explanation is that when we think about the counter-factual claim in a sceptical pressure case, we represent the counter-factual scenario in which the subject's epistemic behaviour falls short of what it would take to count as knowing some proposition. Yet in order to represent the subject in this way, we must also represent the sort of epistemic behaviour that it *would* take for the subject to count as knowing *in that counter-factual scenario*. And for us to represent what it *would take* is precisely to engage in a more complex cognitive strategy than the subject is themselves employing in thinking about the question at hand—reasoning about the quality of evidence, for example. Given our more complex cognitive strategy, we will expect *more* in terms of accrued evidence before making a judgement on the question than the subject we are assessing. Since we *cannot* represent more than one cognitive strategy at a time, the subject's judgment will be assessed relative to our standards for evidence. Consequently, we will deem the subject's evidential basis for their judgement inadequate, and deny that they know.

Might the mechanisms responsible for our negative epistemic judgements in response

to sceptical pressure cases be the same as those which generate the plausibility of sceptical denials? To explore this suggestion, a promising place to start might be in highlighting the strikingly prominent role of counter-factuals in both sceptical pressure cases and sceptical arguments. On the non-sceptical invariantist account of sceptical pressure cases outlined, our denials that the agent knows are the erroneous reflective Type 2 outputs of the dual-process mindreading capacity given the input of counter-factual information. This parallels my working hypothesis concerning sceptical paradoxes: the illusory plausibility of denying we know that sceptical hypotheses are false represents an erroneous reflective Type 2 output of our dual-process mindreading capacity given the input of a sceptical hypothesis. On this basis, I think it is worth exploring whether the non-sceptical invariantist can give an error theory for sceptical paradoxes in terms of dual-process mindreading, in which sceptical hypotheses play the same role as the counter-factual information in sceptical pressure cases.

Many epistemologists think that our appreciation of counter-factual information lies at the heart of the problem posed by sceptical arguments. For example, Michael Williams suggests that an important feature of sceptical hypotheses is that:

the entire course of my experience could be just what it has been even if the world **were** very different from the way I take it to be, or even if there **were** no external world at all. (Williams, 1991, 52, my emphasis.).

As we saw in 1, Pryor too identifies counter-factual considerations as explaining the effectiveness of another form of traditionally threatening sceptical hypothesis:

(1) You are not in a position to know you're not being deceived by an evil demon right now. ... premise (1) is motivated by the thought that no amount of perceptual experience could enable you to determine whether or not you're being deceived by an evil demon, since **you'd** be having exactly the same experiences even if you **were** being so deceived. (Pryor, 2000, 522, my emphasis.).

As we saw in both 2, and again in 4, DeRose diagnoses the coerciveness of the brain in a vat hypothesis in terms similar counter-factual considerations:

[T]he problem with my belief that I'm not a BIV... is that I **would** have this belief (that I'm not a BIV) even if it **were** false (even if I were one). It is this that makes it hard to claim to *know* that I'm not a BIV... For example, and in

particular, if I **were** a BIV, I **would** believe every bit as firmly as I actually do that I *was not* one. (DeRose, 1995, 18-9, my emphasis.).

Finally, we can see the same considerations at play in Pritchard's comments on the same sceptical hypothesis:

I believe that I am not a BIV and that this belief is true. **The problem** with this kind of belief, however, is that if it **were** false—if I **were** a BIV—then I would, by hypothesis, continue to believe that it was true[.] (Pritchard, 2005, 49, my emphasis.).

I think it is clear, from these comments, that the appreciation of counter-factual information concerning sceptical hypotheses plays some role in explaining what is about them that presents us with paradoxes. This lends credence to the idea that the erroneous plausibility of sceptical denials might be explicable in terms of the mechanism giving rise to those erroneous responses to sceptical pressure cases. Our consideration of sceptical hypotheses is often attended by certain counter-factual information. Such counter-factual information requires us to engage in hypothetical thinking, and consequently induces us to more effortful, reflective ways of thinking than we might have otherwise been engaged in. There are certainly clear parallels between sceptical pressure cases and the problematic hypotheses involved in sceptical paradoxes. I think this is a reason to explore my suggestion at the beginning of the section that the sceptical pressure error theory might serve as a good model for an invariantist error theory concerning sceptical paradoxes. In the next section, however, I will review some reasons for thinking that the analogy is not quite as close as it might have seemed.

6.5.4 Dissimilarity of sceptical pressure cases and paradoxes.

In this section, I consider some potential reasons to think that the sceptical pressure error theory described in section 6.5.3 might not serve as a good model for the non-sceptical invariantist to use in explaining away sceptical paradoxes.

The first observation that might raise doubts about this suggestion involves a dissimilarity between sceptical paradoxes and sceptical pressure cases in terms of how the attributer and the subject of the erroneous knowledge judgement are related. The puzzle regarding sceptical pressure cases was that people had apparently inconsistent judgements about what another individual knows between certain cases, which an invariantist maintains do not differ in any way relevant to the truth of the attribution.

People typically judge that an agent knows some proposition p in response to one case, but may also judge that the agent does not know in response to another case which differs only in terms of additional counter-factual information given to the attributer.

In contrast, sceptical paradoxes typically involve an apparently inconsistent set of judgements about what the attributer knows. People typically judge both that they know some contingent proposition p , and that if they know that p , then they know that they are not a brain in a vat, for example, whilst *at the same time* they may judge that they do not know that they are not a brain in a vat. This creates a problem for thinking that the error theory discussed in the previous section can be recruited to adequately explain why we might mistakenly judge that we do not know we are not brains in a vat. This is because the explanation for the erroneous judgement that the subject does not know in response to the sceptical pressure case relies explicitly on the fact that the attributer and subject of attribution are not the *same* person. Nagel sums up the origins of our mistaken epistemic judgement by pointing out that when “one’s own way of thinking about a problem really is more elaborate than that of the *subject* one is observing, one’s intuitive mindreading of this more naive subject may be compromised” (Nagel, 2012, 182). But in case of first-person sceptical paradoxes, it is not clear that the attributer’s way of thinking about the problem *can* be more elaborate than the subject of attribution. This is for the simple reason that they are one and the same person. And for this reason it seems that the error involved in sceptical paradoxes may not be adequately explained in the same way as the error involved in sceptical pressure cases.

The second observation that might raise doubts about the suggestion involves a dissimilarity between sceptical paradoxes and sceptical pressure cases in terms of the content of the apparently conflicting epistemic judgements. The puzzle of sceptical pressure cases, for a non-sceptical invariantist, is that attributers tend to judge that a subject knows that p in one case, and does not know that p in another case. The invariantist’s task is to account for this potentially problematic pattern of judgements. In contrast, sceptical paradoxes are not usually understood as being comprised of a set of judgements that a single subject does and does not know a single proposition. Rather, the potentially problematic pattern of epistemic judgements comprising sceptical paradoxes are those made under a single set of circumstances with respect to a single subject and distinct propositions, such as I know that I have hands, if I know that I have hands, then I know I am not a brain in a vat, and I *do not* know that I am not a brain in a vat.

The crucial dissimilarity here between the puzzles of sceptical pressure cases and sceptical paradoxes concerns the role played by salient ignorance possibilities. These

are possibilities in which some subject *S* is ignorant of some given empirical proposition *p*. In the sceptical pressure puzzle, the possibility that the clock has stopped, and that, subsequently, Wanda does not know that *it is 4:15*, is made salient in case (B). Typically, sceptical paradoxes involve a sceptical possibility in which a subject *S* is ignorant of *every* empirical proposition (that they are a brain in a vat, or dreaming, for example) is made salient in the arguments' first premises. The disanalogy consists in the different ways in which these ignorance possibilities feature in accounting for the puzzles. In the case of sceptical paradoxes, what is to be explained by a non-sceptical invariantist are our judgements about *S*'s epistemic position regarding the salient possibility of their ignorance with respect to every *p*. In contrast, with respect to the sceptical pressure puzzle, what is to be explained are our judgements about *S*'s epistemic position regarding an empirical proposition *p*, given a salient possibility of their ignorance of it.

In the sceptical pressure error theory, the ignorance possibility in which the clock has stopped, was directly implicated in generating the erroneous judgement that the subject did not know the target proposition (ie. what time it is). Even if there is a clear link between hypotheses in sceptical paradoxes and the counter-factual antecedents in sceptical pressure cases, it is also clear that their role in the respective error-theories will not be the same. Simply put, sceptical arguments present us with paradoxes precisely because we normally *continue* to judge that we know various empirical propositions, even if considerations of sceptical hypotheses compel us to judge that we do not know we are not brains in vats. And for this reason too it seems that the error involved in sceptical paradoxes may not be adequately explained in the same way as the error involved in sceptical pressure cases.

There is one final observation that I think may cast doubt on any close analogy between the error theory described above with respect to sceptical pressure cases, and one applicable to sceptical paradoxes. This concerns a significant difference between the sorts of counter-factual considerations involved. Central to the explanation of the erroneous judgements that the subject did not know in sceptical pressure cases was the fact that the counter-factuals required us to envisage the subject of attribution collecting more evidence than they in fact do. But this is not true of the counter-factual considerations that usually accompany radical sceptical hypotheses. For example, in order to understand the counter-factual 'if I were a brain in a vat, then I would not be able to tell that I was not one', I may have to envisage being a brain in a vat. But at no point in doing so, do I plausibly envisage any further evidence that I can acquire, whether in actuality, or in the counter-factual scenario. Rather, it is characteristic of

radical sceptical hypotheses such as this one that there simply is no more evidence I can acquire in actuality with respect to any contingent proposition, and especially not in the scenario described by the hypothesis. Nagel's account of the erroneous judgements, that the subject in sceptical pressure cases did not know the target proposition, was based on the idea that "understanding the road not taken—the hypothetical collection of surplus information—requires the representation or replication of a more elaborate cognitive strategy" (Nagel, 2012, 184) than the subject of attribution *actually* employs. With radical sceptical hypotheses and their related counter-factual considerations, however, there is no hypothetical collection of surplus information. And for this reason, again, it seems that the sort of error involved in sceptical paradoxes may not be adequately explained in the same way as that in sceptical pressure cases.

I have outlined some reasons for thinking that Nagel's approach to explaining away the pattern of judgements about knowledge ascriptions concerning sceptical pressure cases might not serve as an analogous model for providing a non-sceptical invariantist error theory for sceptical paradoxes. The question of whether these points count decisively against the idea that the erroneous judgements involved in sceptical paradoxes are explicable in the same way as those in sceptical pressure cases would be a subject for further work. In the remainder of this chapter, I will move on to consider a promising model for a non-sceptical invariantist error theory, that further develops my driving thought that it is a kind of 'internal conflict' that gives rise to sceptical paradoxes.

6.6 Dual Processes and Default-Interventionism.

I have surveyed two potential routes for an invariantist to provide an adequate error theory with respect to sceptical paradoxes. The first proposed that the availability heuristic could explain why we might mistakenly judge that we do not know sceptical hypotheses are false, and it was argued to face significant issues. The second proposed that an account of why we might mistakenly judge that a subject does not know a proposition in response to sceptical pressure cases might serve as a potential model for the apparent plausibility of sceptical arguments' first premises. I suggested that this model too faces some significant issues if it is applied to sceptical paradoxes. In what follows, I put forward third potential model for an invariantist error theory concerning sceptical paradoxes. This model is inspired by another way which cognitive psychology has been used to provide an invariantist account of some potentially problematic patterns of epistemic judgement.

Gerken (2012, 2013) has presented “an account of the cognitive processes involved in the formation of judgements about knowledge ascriptions [that] may provide a unified explanation of” (2013, 2) certain data concerning knowledge ascription which may motivate rivals to non-sceptical strict invariantism. Gerken’s work informs my aims, and specifically, it will be worth explaining the bearing of the following comments on my work:

“My working assumption that judgements about knowledge ascriptions do not differ radically from other judgements involves the assumption that such judgements are also constrained by our cognitive capacities. They may, therefore, be biased. I have applied this general idea to patterns of judgements about knowledge ascriptions in a manner that is compatible with non-skeptical strict invariantism.” (Gerken, 2012, 162)

Like Gerken, my aim here is to give an non-sceptical strict invariantist account of a potentially problematic pattern of judgements about knowledge, in terms of the dual process theories of the cognitive capacities underlying these judgements. Specifically, I aim to describe a non-sceptical strict invariantist error theory concerning sceptical paradoxes, in terms of the dual process theories of judgement. In developing this error theory I draw upon the *default-interventionist* version of dual-process theory defended by Evans (2007; 2013; 2011; 2014). In the next and final section, I draw upon the default-interventionist framework, as well as the results from chapters 2 and 3, in order to present a new and promising error theory for non-sceptical invariantists. First, in the remainder of this section, I outline the default-interventionist account of dual process theory.

Default-interventionism is a refinement of dual process theory of higher cognition which describes the way in which Type 1 and Type 2 processes are employed and interact in cognitive tasks.¹⁰ The most basic claim of default-interventionism is that the processing involved in a given cognitive task proceeds in the following way: “Type 1 processing produces a rapid and intuitive default response, which may or may not be intervened upon by subsequent Type 2 reasoning which is slower and deliberative in nature” (Evans, 2011, 93). Or, in more detail:

Type 2 processing requires extremely limited and precious working memory resources, according to the definitions we are putting forward. These

¹⁰ However, there are dual process theories that are not default-interventionist in their account of how processes interact in cognitive tasks. See, for example, (Sloman, 1996), and (Smith and DeCoster, 2000).

must be selectively allocated to the most important task at hand. Default interventionism allows that most of our behavior is controlled by Type 1 processes running in the background. Thus, most behavior will accord with defaults, and intervention will occur only when difficulty, novelty, and motivation combine to command the resources of working memory. (Evans and Stanovich, 2013, 238)

Perhaps the most fundamental aspect of default-interventionist versions of dual-process theory is, as the above quote clearly states, that humans—being the finite beings they are—are simply not capable of engaging in reflective thought to undertake every cognitive task that might come our way. Most of the time, Type 1 processes carry the majority of the cognitive load, with Type 2 processes only being called upon to do only as much as is absolutely necessary. Importantly, and in contrast to more classical accounts, Type 1 and Type 2 processes do not occur in isolation, but rather *interact* with each other frequently. And it is this interaction that default-interventionism purports to explain. To see how it does, consider Figure 1 below, which models the interaction of Type 1 and Type 2 processes on a recently developed default-interventionist account of dual process theory.

The diagram in Figure 1 illustrates the default-interventionist model of higher cognition. As noted, each cognitive task prompts an autonomous Type 1 process which generates a response, here labelled ‘A1’. In explaining this aspect of the model, Evans states that “A1 simply represents the default intuitive response that will be made unless Type 2 processing intervenes with rethinking the problem” (Evans, 2011, 95). The box labelled ‘Reflective processing (Type 2)’ in the diagram represents the engagement of Type 2 processing. On default-interventionist theories, such reflective processing always occurs subsequently to Type 1 processing to take up and assess the intuitive responses generated by it.¹¹ The actual Type 2 assessment of the intuitive response A1 is represented by the decision box in the diagram labelled ‘A1 justified?’. On the default-interventionist theory there are two possible outcomes when this reflective “attempt is made to verify the initial intuition A1” (Evans, 2011, 94): either the Type 2 process accepts the Type 1 response A1, or else it rejects it. The former possibility results in the reflective endorsement of A1; it becomes the ‘default’ response. The latter possibility results in an attempt to rethink the problem reflectively; this is what is termed ‘intervention’, or ‘override’.

¹¹ See (Evans, 2007) and (Faghihi et al., 2015).

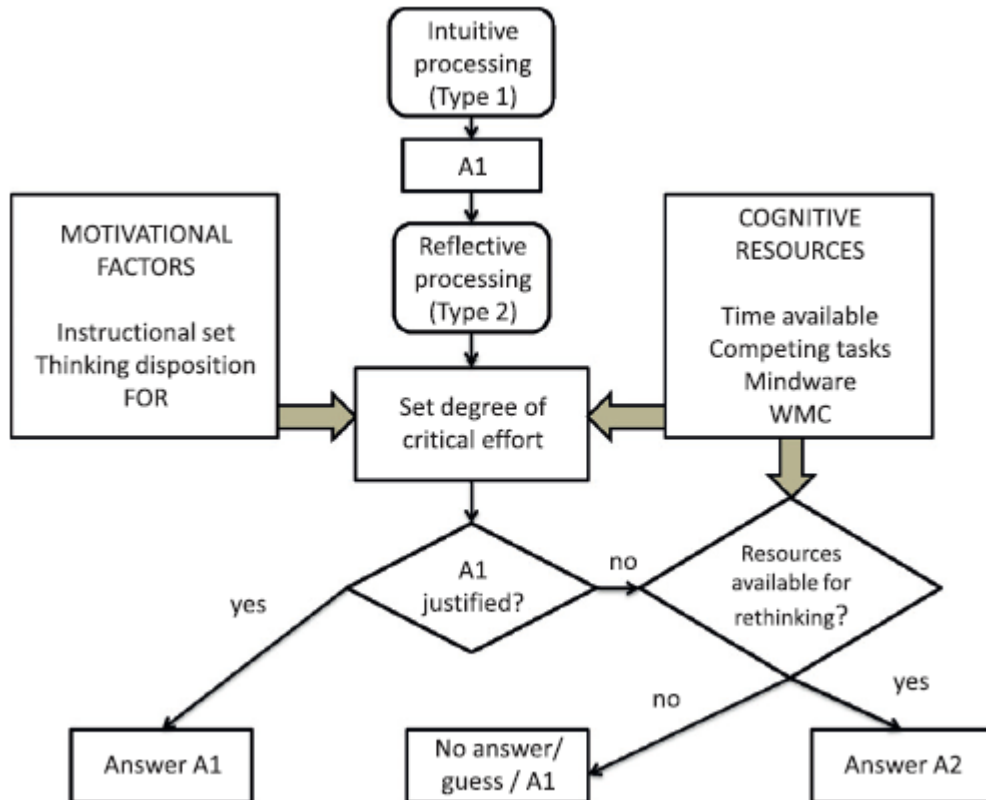


Figure 1: Default-Interventionist model of Dual Process Theory.¹²

The propensity to accept or reject an initial intuitive Type 1 response is determined by a degree of critical effort “which is itself a function of the motivational and cognitive factors” (Evans, 2011, 96). The motivational factors include the Feeling of Rightness (or FOR), any instructions accompanying the task (or *instructional set*), and *thinking dispositions*. The Feeling of Rightness is the term for a meta-cognitive judgement of a varying strength determined by “the fluency with which Type 1 processes produce an initial answer” and “that in turn mediates the extent of Type 2 engagement” (Thompson et al., 2013, 238).¹³ An instructional set is basically the set of instructions that an individual is given for completing a specific task.¹⁴ Thinking dispositions refer to the dispositions different individuals may have for thinking about problems, for example, some people may tend to approach problems with analytic thinking, whilst others may tend to use more intuitive approaches. The cognitive factors include time constraints on a task, competing tasks include additional tasks that involve working memory use,

¹² From (Evans, 2011).

¹³ See also (Thompson et al., 2011; Thompson and Johnson, 2014).

¹⁴ See also (Macpherson and Stanovich, 2007), Cf. (Auf, 2009) and (Set, 2009).

whilst working memory capacity is self-explanatory. Finally, ‘mindware’ refers to the “rules, procedures, and strategies that can be retrieved by the algorithmic and reflective minds” and is crucial for “the construction of an alternative response to substitute during the override of Type 1 processing” (Stanovich et al., 2011, 107).

The resources of default-interventionism capture the notion of an internal conflict, and can be drawn upon to further develop my idea for a general strategy involved in resolving sceptical paradoxes in terms of a psychological error theory. This strategy views the internal conflict involved in sceptical paradoxes as instances of what are called *cognitive illusions*. In the most basic sense, cognitive illusions are systematically generated false beliefs, which are “just the by-products of a cognitive system which is responsible for extracting knowledge about the world” (Yarritu et al., 2015, 2). Cognitive illusions have been described as being “analogous to optical illusions in leading to errors we commit *without knowing we are doing so*, except they arise from our difficulties in quantifying and dealing with probabilities, uncertainty, and risk” (Nicholls, 1999, 1386). Cognitive illusions can be seen in the responses to various problems which form the Cognitive Reflection Test (CRT).¹⁵ In particular, the item known as the ‘bat and the ball’ problem has become a well-known case of a cognitive illusion, and serves as an instructive example of how default-interventionism accounts for certain internal conflict. In the bat and the ball item of the CRT, participants are tasked with solving the following problem:¹⁶

Bat and ball: A bat and a ball cost \$1.10 in total. The bat costs \$1 more than the ball. How much does the ball cost?

The bat and ball problem presents a cognitive illusion since the intuitive answer to the problem is normally judged to be \$0.10. This answer is incorrect (it is \$0.05), but yet it is a mistaken response that was given more than half of the time by participants from Ivy league universities (Evans, 2011). The felt *plausibility* of the erroneous answer “\$0.10” is the cognitive illusion here. When faced with the problem, a Type 1 process generates an initial answer cued by contextual features of the problem. The intuitive answer ‘\$0.10’, for example, may be cued by the salience of the information that the bat costs \$1 more than the ball, and that together they cost \$1.10, which may prompt a Type 1 process to preconsciously subtract \$1.00 from \$1.10. This erroneous answer then may be the default response to the problem if either Type 2 processes are not

¹⁵ See (Frederick, 2005).

¹⁶ From (Toplak et al., 2014).

engaged in critical assessment of it, or if reasons are not found for this Type 1 output to be overridden and replaced by Type 2 response.

The notion of a cognitive illusion can be related to the explanandum of a non-sceptical invariantist error theory concerning sceptical paradoxes, viz. that we find it plausible to judge that we don't know that sceptical hypotheses are false. The non-sceptical invariantist can explain our tendency to be systematically misled by sceptical arguments in this way as an example of a cognitive illusion. In providing their error theory concerning our judgements with respect to sceptical premises, they can maintain that these are the result of facts about our cognitive architecture. In the example of the bat and the ball problem, the felt plausibility of the erroneous answer '\$0.10' may be accounted for on a dual-process theory picture as the result of our cognitive architecture in terms of the interaction of Type 1 and Type 2 processes. In the case of sceptical paradoxes, perhaps the non-sceptical invariantist can now use dual-process theories to account for the felt plausibility of erroneously denying that we know that we are not in sceptical scenarios. On this proposal, the internal conflict of sceptical paradoxes is accounted for a cognitive illusion. Specifically, the felt plausibility of the first premises of sceptical arguments is viewed as the erroneous result of Type 2 processes. On this picture, the sceptical paradoxes represent cognitive illusions, in the same sense that the Bat and the Ball item, but which differ in that the illusory judgement is a Type 2, rather than a Type 1 output.

Before further developing this account, it will be useful to use another example to draw out the idea of a cognitive illusion in which it is a Type 2 judgment that is erroneous. To do this, let's explore how the default-interventionism picture could be used to account for the erroneous knowledge ascriptions made in response to the sceptical pressure case discussed by Nagel. To recap, the non-sceptical strict invariantist maintains that an attribution of knowledge to Wanda in case (A) represents a true positive. In contrast, they will maintain that a denial that Wanda knows in case (B) represents a false negative.

In the case of (A), following the default-interventionist picture, the invariantist can account for the judgement that Wanda knows in terms of the *default* aspect of our cognitive architecture. On this account, the task facing the reader of the vignette prompts a Type 1 process to generate an answer. An intuitive answer that Wanda knows in this case is strongly cued by contextual features of the vignette, such as the subject's behaviour and their environment. Intuitively, we regard Wanda as knowing that it is 4:15 on the basis of her coming to belief this having checked a clock which displays the actual time of 4:15. Now, using the default-interventionist picture, the invariantist can suggest that in the absence of any obvious reason to reflect on the output of the Type 1

process with costly Type 2 processing, the intuitive answer that Wanda knows that it is 4:15 will be given by default.¹⁷ A non-sceptical invariantist typically maintain that Wanda does know in (A), and so can regard this intuitive answer as a true positive.

In the case of (B) the invariantist can draw on the default-interventionist picture to account for the judgement that Wanda does not know in terms of both the *intervention*, and subsequent *override* aspects of our cognitive architecture. Here, the account proceeds initially in the same manner as with case (A). Reading the vignette (B) prompts a Type 1 process to generate an answer cued by contextual features of the case. In contrast to (A), however, the counter-factual content involved in (B) prompts a Type 2 process to intervene upon the output of the initial Type 1 process. This occurs because, as noted above, Type 2 processing “is involved whenever we engage in hypothetical thinking, supposition, or mental simulation” (Evans, 2007, 110). The intervention of a Type 2 process involves the subsequent critical examination of the intuitively generated answer. On the default-interventionist picture, if the product of an initial Type 1 process is reflectively deemed unsatisfactory, it will be overridden and the reflective answer that Wanda does not know is given instead. A non-sceptical invariantist typically maintains that Wanda does know in (B), and so will regard this reflective answer as a false negative.

We have seen that the resources of default-interventionist theories provide a way of characterising internal conflicts as cognitive illusions. In this way, the non-sceptical invariantist can draw on these resources to account for the internal conflict of sceptical paradoxes as cognitive illusions. In contrast to the cognitive illusions discussed in the psychology literature, however, my proposal is that the erroneous judgment in sceptical paradoxes represents the output of a Type 2 process. To explore this idea, I applied the resources of default-interventionism to the example of sceptical pressure cases discussed by Nagel. I suggested that it provides a promising account of sceptical pressure cases as examples of cognitive illusions involving erroneous Type 2 judgements. Ultimately, however, I think that various salient dissimilarities between sceptical paradoxes and sceptical pressure cases, put a limit on how closely an invariantist can adequately model an error-theory for the former upon the latter.

To summarise this section, and the chapter so far, I think that two important points have been made concerning the prospects of providing a happy-face non-sceptical invariantist resolution of sceptical paradoxes. Firstly, internal conflicts such as sceptical paradoxes, can be account for as cognitive illusions arising from our cognitive architec-

¹⁷ What is significant here is that, in this case, is that a Type 1 response is given to the case *without* the intervention of a Type 2 process. It is an important feature of the default-interventionist picture that Type 2 processes do not necessarily have to be involved in judgement making.

ture involving Type 1 and Type 2 processes. Second, some cognitive illusions can be attributed to Type 2 processes being responsible for generating erroneous judgements—as in the case of sceptical pressure cases. Interestingly, this is a commitment of default-interventionist dual-process theory that is commonly over-looked:

Perhaps the most persistent fallacy in the perception of dual-process theories is the idea that Type 1 processes (intuitive, heuristic) are responsible for all bad thinking and that Type 2 processes (reflective, analytic) necessarily lead to correct responses. (Evans and Stanovich, 2013, 229)

This point is, I think, especially significant in that it raises the interesting and novel possibility that sceptical paradoxes arise because of an error attributable to reflective Type 2 thinking. In the next section, I take on this idea, and apply the default-interventionist picture to the case of sceptical arguments to describe a more detailed non-sceptical invariantist account of sceptical paradoxes as Type 2 cognitive illusions.

6.7 Default-interventionism, Sceptical Illusions, and Invariantism.

In this section, I present a non-sceptical strict invariantist error theory concerning sceptical paradoxes by drawing together the results of two earlier chapters with the default-interventionist model of dual process theory. I begin by briefly recapping the main points of the chapter concerning what was identified as required of an invariantist solution to sceptical paradoxes, and what was found. Next, I set out some general ways in which the default-interventionist model of dual-process theory can be utilised by an invariantist in an error theory to resolve the sceptical paradox. Finally, I draw on the findings of chapters 2 and 3 concerning, respectively, sceptical hypotheses and sceptical arguments, together with the default-interventionist framework to provide the details of the proposed invariantist error theory.

At the beginning of this chapter, I described what was required of an invariantist to satisfactorily resolve the sceptical paradoxes. This amounts to doing two things. The first is to identify the false claim amongst the set of the two premises, and the negated conclusion of a sceptical argument. The second is to provide an adequate error theory concerning the claim that was identified as false, viz. a plausible account of why we might have mistakenly judged it to be true. A non-sceptical strict invariantist will maintain that

we do know the negations of sceptical hypotheses, and so identify the minor premise of a sceptical argument as false. I have suggested that a promising route for an invariantist would be to provide a psychological, rather than a semantic error theory. So let me offer such an account in what follows drawing on my favoured default-interventionist picture.

Having set out the default-interventionist framework in the previous section, it now remains for me to apply it to sceptical paradoxes, and begin to describe my proposal for an adequate psychological error theory. I set out this proposal in two stages. First, I set out some ways in which the invariantist can directly apply the framework to the sceptical problem to give a psychological error theory. Here, the main work involves characterising the internal conflict of sceptical paradoxes as a cognitive illusion, in terms of the default-interventionist resources. In the second stage, I draw upon my account of sceptical hypotheses, and epistemic principles to suggest an account of the mechanism responsible for the hypothesised intervention, and subsequent override.

As I have already noted, a non-sceptical strict invariantist will maintain that the first premise of a sceptical argument is false. They need to explain why we might judge it to be true. On my account, the invariantist can use the default-interventionist picture to characterise a judgement that we do not know sceptical hypotheses to be false as an erroneous output of our cognitive processes. The first step on my account is to identify the erroneous judgement—I do not know I am not a brain in a vat, for example—as the output of a reflective Type 2 process, rather than an intuitive Type 1 process. In other words, on my non-sceptical invariantist account, sceptical paradoxes arise because our reflective judgments concerning minor sceptical premises conflict with our intuitive knowledge ascriptions. When confronted with sceptical arguments, we may feel a reflective pull towards denying we know that we are not brains in vats, for example, even whilst we intuitively think we know we have hands. The internal conflict lies between the outputs of our Type 1 and Type 2 cognition, just as it did in the Bat and the Ball example of a cognitive illusion. The difference is that in the latter case, the illusion was the intuitive plausibility of the erroneous Type 1 response ‘\$0.10’. In the case of sceptical paradoxes, on my account, the illusion is the reflective plausibility of the Type 2 judgements concerning sceptical minor premises, such as ‘I do not know I am not a brain in a vat’.

It is worth pointing out that the assumption that our erroneous sceptical judgements as reflective Type 2 judgements finds support in the literature on sceptical paradoxes. Epistemologists commonly recognise that finding the minor premises of sceptical arguments to be plausible is something that can happen only *on reflection*. Williams, for example, thinks this, even if he confusingly uses his own distinct technical sense of the

term ‘intuitive’, as can be seen in the following comments:

However damaging the sceptic’s conclusions may be, they will amount to a threat only if we suspect that we are committed to them. Thus the sceptic’s arguments must at least appear to be compelling because intuitive. By calling them “intuitive”, I mean that the sceptic must derive his conclusions from what anyone one can see, *on reflection*... (Williams, 2012, 354, my emphasis)

skepticism is not a stable, long-term option, it can be attained, or at least approached, in the context of *sustained philosophical reflection*, which detaches us from the perceptual and practical factors that ordinarily influence belief. (Williams, 2013, 40)

This plausible assumption, that our erroneous judgements with respect to the minor premises of sceptical arguments are the outputs of Type 2 reflective processes, has some further implications for the proposed account. The default-interventionist model of dual-process theory holds that, in forming a judgement “fast Type 1 processing generates intuitive default responses on which subsequent reflective Type 2 processing may or may not intervene” (Evans, 2011, 213). The next step in presenting my non-sceptical invariantist account, then, is to suggest that sceptical paradoxes arise because of *override*. On my view, an erroneous categorical judgement that a sceptical hypothesis is *not* known to be false is the result of Type 2 processes overriding a default Type 1 intuitive judgement, and replacing it with a Type 2 reflective judgement.

I have argued that default interventionist dual-process theories of judgement can help a non-sceptical invariantist account for the erroneous plausibility of minor sceptical premises. This is the result of a Type 2 reflective process intervening upon, and overriding an initial Type 1 process. To complete the proposed invariantist psychological error theory concerning sceptical paradoxes, it remains to give an account of *what* triggers this reflective intervention, and *why* it results override. In the remainder of this section, I set out an account of both the mechanism for Type 2 intervention, and of the subsequent override. Here, I draw on the further resources of epistemology, and specifically my accounts of sceptical hypotheses and sceptical arguments from chapters 2 and 3, respectively. I combine these accounts with default-interventionist resources to suggest how sceptical paradoxes arise as the result of reflective override.

What mechanism triggers the intervention of Type 2 processes when considering sceptical arguments? Following the suggestion earlier, I suggest that the negation involved in the minor sceptical premise prompts a Type 2 process to intervene. Recall that

these premises deny that we know that sceptical hypotheses do *not* obtain. Drawing on the above discussion of sceptical pressure cases, I propose that this negation can be identified the mechanism triggering Type 2 intervention. In order to think about the negated possibility (that I am a brain in a vat, for example), it is necessary to sustain a decoupled representation, and this requires Type 2 processing:

[I]n order to reason hypothetically, we must be able to prevent our representations of the real world from becoming confused with representations of imaginary situations. The so-called cognitive decoupling operations are the central feature of Type 2 processing that makes this possible. (Evans and Stanovich, 2013, 236)

The first part of the account, then, maintains that when we are confronted with sceptical arguments, the negation of sceptical hypotheses triggers Type 2 processes to intervene, and, subsequently, generate our reflective judgements about the arguments' premises.

The next thing to spell out, on this account, is why the Type 2 process intervention triggered by minor sceptical premises results in override. In other words, why might our Type 2 processes generate the reflective plausibility of premises such as 'I do not know I am not a brain in a vat'? My suggestion is that this override occurs because a reflectively available adequacy constraint on responses to sceptical arguments' first premises can never be satisfied. On my view, the epistemic principles underwriting arguments from ignorance combine with the characteristic generality of sceptical hypotheses to create an unsatisfiable constraint on rejecting minor sceptical premises. In short, any counter a claim of the form 'S knows that not-*sh*', when *sh* is a sceptical hypothesis with the widest scope, will always seem to beg the question.

This aspect of my proposed account can be now described in more detail. It will be instructive to focus on the example of a paradoxical sceptical argument, and so I will use the argument presented in section 6.2—essentially ENVATTED BRAIN—as a toy model here. How does this argument generate a cognitive illusion in the sense outlined above. In proceeding, I assume that when someone is presented with the first premise of a sceptical argument (for example, I do not know that I am not a brain in a vat), the negation of the brain in a vat hypothesis prompts a Type 2 process that intervenes upon any Type 1 process. Now, my hypothesis is that my Type 2 process will critically assess the plausibility of the sceptical premise in terms of the principle of epistemic possibility that I argued for in chapter 3:

Epistemic Possibility.

For all S, p , S knows that p only if it is not epistemically possible for S that $\neg p$.

Here, then, is a suggestive sketch of what sort of processes might give rise to the illusory, reflective judgement that (1) is plausible. Suppose I encounter the sceptical argument presented in 6.3. With the *Epistemic Possibility* principle in mind, I might reflect that I know I am not a brain in a vat only if it is not epistemically possible for me that I am a brain in a vat. I might also reflectively recognise, given the analysis of epistemic possibility that I favour, that it is not epistemically possible for me that I am a brain in a vat only if my being a brain in a vat is inconsistent with me knowing what I know. When reflecting on the sceptical premise that ‘I do not know that I am not a brain in a vat’, I may realise, then, that this premise is false only if, given what I *know*, it *cannot* be the case that I am a brain in a vat. On my view, the reflective plausibility of the sceptical premise that ‘I do not know I am not a brain in a vat’ is the illusory product of Type 2 intervention and override. My suggestion is that, insofar as I feel a reflective pull towards assenting to this premise, this is merely the result of Type 2 override. What I propose is that, when reflecting, I may be led inexorably to judge that I do not know I am a brain in a vat, given my appreciation of the *Epistemic possibility* principle, and my inability to satisfactorily counter the claim that it *is* epistemically possible for me that I am a brain in a vat. In the remainder of this section, I set out this proposal in more detail.

It is worth stressing, however, that I do not take that the view that *Epistemic Possibility* ought to be rejected by a non-sceptical invariantist. The role played by this principle in my account of Type 2 override in sceptical paradoxes is subtle. On my account, it is not that I infer that I do not know I am not a brain in a vat, from *Epistemic Possibility* and the *de facto* recognition that it is epistemically possible for me that I am one. Rather, the pull I feel, *on reflection*, towards the sceptic’s minor premise arises because of two things: my recognition that *Epistemic Possibility* holds, and the *de facto* impossibility of *satisfactorily* maintaining that it is not epistemically possible for me that I am a brain in a vat.

Why might that be? After all, various examples of propositions I know, such that it could not be the case that I am a brain in a vat *and* know these propositions might come to my mind readily. I could think, for example, “I have stood for many hours in my office, I know *that*, and so I cannot be a brain in vat.” This is clearly a valid counter to the claim that it is epistemically possible for me that I am a brain in a vat, for example.

This is obviously not, however, a satisfactory way to counter this claim in the context of a sceptical argument. This is because my counter-example to the claim that the sceptical hypothesis is epistemically possible for me contravenes the readily appreciable injunction not to beg the question against the argument that I am considering. But this sceptical argument I am considering threatens to challenge all my putative empirical knowledge at once, and my presumed knowledge of having stood for many hours in my office is, of course, a piece of empirical knowledge. I clearly beg the question against the sceptical argument, then, if I attempt to reject the premise that I do not know that I am not a brain in a vat by citing my putative knowledge that I have stood for many hours in my office.

Worse still, as I reflect further, I realise that I can never be in a position to *satisfactorily* deny that it is epistemically possible for me that I am a brain in a vat. Any attempt I might make to do this is determined to fail. This is because, in order to do so I would have to consider myself to have ruled out that I am a brain in a vat, *without* begging the question against the sceptical argument. I can rule out that I am a brain in a vat only if this hypothesis is incompatible with what I know. And that, in turn, would require me to know some empirical proposition. Yet, as I argued in chapter 2, sceptical arguments threaten to challenge all empirical knowledge in virtue of sceptical hypotheses being possibilities of *global* empirical ignorance. Consequently, I beg the question against the sceptical argument by relying on *any* putative empirical knowledge of mine.

At this point, the result of my reflections in considering the sceptical argument might be that I have no way to adequately reject the premise that I do not know that I am not a brain in a vat. I may find myself reflecting along the following lines: I feel like I really do know that I am not a brain in a vat, but, on reflection, there is just no way to deny that I do not know that, so what can I do? In this way, it might be that I begin to feel that the sceptical premise, that I do not know that I am not a brain in a vat, has a sort of reflective, albeit *peculiar and illusory* plausibility. On this picture, the illusory plausibility of sceptical premises represents a sort of *cognitive attrition*—a feeling of “if you can’t beat them, join them”—and can be characterised by the non-sceptical invariantist as an example of override, and a *cognitive illusion*. And this, I think, neatly captures the sense in which sceptical paradoxes represent an internal conflict; we do not really *accept* that we do not know that we are not brains in vats, but we also recognise that we are not able to reflectively *reject* this claim, given the natural considerations of *epistemic possibility* involved.

According to default-interventionist theories, costly Type 2 processing cannot continue

indefinitely, and must end at some point. On my account, the intervention of Type 2 processes terminates in override. This is because the injunction not to beg the question, and the characteristic generality of sceptical hypotheses together imply that I cannot counter the claim that it is epistemically possible for me that I am a brain in a vat. My intervening Type 2 process must inevitably end, and it will do so without having satisfactorily modeled the falsity of the sceptical premise that I do not know that I am not a brain in a vat. In the context of considering sceptical arguments, the impossibility of adequately rejecting a sceptical premise *reflectively* results in override, and gives rise to an illusory reflective judgement that it is true.

In this way, I suggest, the non-sceptical invariantist can account for the erroneous plausibility of sceptical premises, and, subsequently, why sceptical paradoxes emerge. The plausibility of a sceptical argument's minor premise is illusory, and a sceptical paradox is a cognitive illusion.

6.8 Assessing my proposal.

In the previous section, I outlined a non-sceptical invariantist error theory concerning sceptical paradoxes based on a default-interventionist model of dual-process theory of cognition. I turn now, in this final section, to comment on how this error theory relates to the residual issues of previous chapters. In chapter 1, I considered various constraints on a fully satisfactory, happy-face resolution of sceptical paradoxes. The most important of these was that a solution must provide an error theory concerning the set of judgments comprising the paradoxes. I further argued that an important constraint on the adequacy of these error-theories is that the explanation for the erroneous judgement must be given in terms that are not purely *epistemological*.

On my account, the explanation for the erroneous pull of sceptical premises is given in terms of dual-process theories of cognition. Our systematic error when presented with sceptical arguments can be described as the result of our cognitive architecture. A mistaken judgement that we don't know that sceptical hypotheses are false amounts to a Type 2 process output given the cognitive task of assessing sceptical arguments. This happens because Type 2 processes are triggered and override our initial automatic attempts to generate a Type 1 judgement about sceptical premises. Yet Type 2 processes will tend to generate negative judgements that sceptical hypotheses are not known to be false. This happens, on my view, as a result of the reflective appreciation of the *Epistemic Possibility* principle, and the impossibility of generating a non-question-begging counter

to the claim that the sceptical hypothesis is epistemically possible. Type 2 processing can not continue indefinitely, and must stop somewhere. The result is that our cognitive efforts end with effective override by Type 2 processes, and we feel a reflective pull towards denying we know that sceptical hypotheses are false.

In this way, I think that I have succeeded in sketching an error-theory for a happy-face resolution of sceptical paradoxes that is broadly confusion-therapeutic. I did not, in presenting my error theory, make any obvious conceptually concessive, or revisionary, moves to resolve the sceptical paradoxes. The epistemic principle I argued to underwrite sceptical arguments in 3 need not be *conceded* to the sceptic, or given up because of its roll in the explanation for how sceptical arguments work. Rather, the essential point of the *cognitive* error theory that I have put forward is that *Epistemic Possibility* is a *bona fide* and powerful part of our ordinary ways of thinking about knowledge. The fact that sceptical arguments can confuse us, and are best understood as underwritten by *Epistemic Possibility*, neither implies that this principle is extraneous to our concept, nor that it should be given up. On the account of sceptical paradoxes as *cognitive illusions* that I have proposed, there is no ‘glitch’ (to use Schiffer’s phrase) in our concept of knowledge. Rather, a glitch in our cognitive architecture, given inputs such as sceptical hypotheses, gives rise to illusion that the minor premises of sceptical arguments—let alone their conclusions!—have any *genuine* plausibility based on *Epistemic Possibility* at all.

Another residual worry about this default-interventionist dual-process account is whether or not it is equally compatible with both sceptical and non-sceptical invariantism. Again, here, I draw on an insight from (Nagel, 2011) to suggest that this is not the case. She has pointed out that dual-process accounts of certain patterns of epistemic judgements cannot be considered to be equally friendly to the sceptical invariantist who denies we know any empirical propositions. In particular, if the default-interventionist account of sceptical paradoxes that I have outlined here is correct, then, “the skeptic cannot claim support from this pattern of [judgements]...without an independent argument to establish that skeptical controlled [ie. reflective Type 2] judgments are epistemically superior to non-sceptical controlled and to automatic [ie. intuitive Type 1] judgments” (Nagel, 2011, 22). The point here is that the best—and perhaps, only— motivation for scepticism is precisely the undeniably felt plausibility of their premises. Insofar as the default-interventionist account identifies this plausibility as merely illusory, then the “motivation for skepticism is undercut” (ibid) by this non-sceptical invariantist error-theory. An error-theory for sceptical paradoxes, that offers “a properly detailed diagnosis

and expose of its power to seduce” (Wright, 1991, 89). And that, I take it, is just what a happy-face resolution to sceptical paradoxes should do.

In the previous chapter, I argued that the semantic blindness objection counted against various proposed contextualist resolutions to sceptical paradoxes. The question remained, however, whether this meant that a non-sceptical invariantist could fare any better on this scorecard by providing an error theory that had any advantages to the succumb to similar problems. I think that the non-sceptical invariantist psychological error theory I have outlined does have significant advantages over the contextualist’s semantic error theory. The main problem that afflicted the semantic error theory was that semantic blindness is implausible for recognised context-sensitive terms. Though not a knock-down objection—the contextualist is free to suggest that ‘knows’ exhibits a novel sort of context-sensitivity—this raises some significant issues with respect to error theory. This semantic error theory purportedly explains our mistaken judgement regarding sceptical arguments only in terms of a further mystery: why does the context-sensitivity of ‘knows’ give rise to semantic blindness at all. An error theory that tries to explain some phenomenon only by pointing to something that itself is mysterious, does not really explain much at all.

In contrast, the invariantist psychology error theory I have described does not suffer from the same secondary explanatory issues as the contextualist’s semantic error theory. On this view, the sort of error we make when faced with sceptical arguments is ubiquitous in human reasoning. Unlike semantic blindness, there is no obvious mystery involved in the proposed psychological error theory. It explained the mistaken judgement partly comprising the sceptical paradox, in terms of independently motivated epistemological analyses of sceptical arguments, and independently motivated theories about the architecture of human cognition. In this way, the invariantist enjoys an advantage over the contextualist on the scorecard when it comes to providing an error theory to satisfactorily resolve sceptical paradoxes. Moreover, it is not clear that the contextualist can help themselves to the psychological error theory. Contextualists typically take the plausibility of sceptical premises as evidence for the context-sensitivity of ‘knows’; we find them plausible because we recognise that in some contexts they express truths. Yet the psychological error theory I propose maintains that this plausibility is a cognitive illusion. So adopting this account would undermine the motivation for contextualism that sceptical paradoxes might offer.

6.9 Concluding remarks.

In this chapter, I considered various potential ways of utilising the resources of cognitive psychology to provide a non-sceptical invariantist error theory concerning sceptical paradoxes have been critically assessed. I argued that explanations for the why the sceptical paradox arises by accounting for our erroneous judgement in terms of the availability heuristic failed to be adequately explanatory. Next, I considered a suggested explanation that appealed to a dual-process theory of the mindreading capacity, and argued that it too failed to be adequately explanatory. Finally, I proposed a new invariantist error theory based on the framework of a default-interventionist dual process theory of judgement, and the findings of earlier chapters concerning sceptical hypotheses and arguments. I conclude that the *cognitive* error theory that I have described here could indeed represent a invariantist happy-face resolution of sceptical paradoxes.

BIBLIOGRAPHY

- (2009), "Aufgabe." In *The Penguin dictionary of psychology* (A. Reber, R. Allen, and E. Reber, eds.), London: Penguin, URL <http://search.credoreference.com/content/entry/penguinpsyc/aufgabe/0>.
- (2009), "Set." In *The Penguin dictionary of psychology* (A. Reber, R. Allen, and E. Reber, eds.), London: Penguin, URL <http://search.credoreference.com/content/entry/penguinpsyc/set/0>.
- Almeida, Claudio (2012), "Epistemic closure, skepticism and defeasibility." *Synthese*, 188, 197–215.
- Anderson, Charity (2014), "Fallibilism and the flexibility of epistemic modals." *Philosophical Studies*, 167, 597–606.
- Baker, Carl (2012), "Indexical contextualism and the challenges from disagreement." *Philosophical Studies*, 157, 107–123.
- Baumann, Peter (2010), "The case for contextualism: Knowledge, skepticism, and context, vol. i – keith derose." *Philosophical Quarterly*, 60, 424–427.
- Beebe, James (2010), "Constraints on sceptical hypotheses." *Philosophical Quarterly*, 60, 449–470.
- Blaauw, Martijn (2005), "Challenging contextualism." *Grazer Philosophische Studien*, 69, 127–146.
- Blome-Tillmann, Michael (2008), "The indexicality of 'knowledge'." *Philosophical Studies*, 138, 29–53.
- Blome-Tillmann, Michael (2009), "Knowledge and presuppositions." *Mind*, 118, 241–294.
- Boult, Cameron (2013), "Epistemic principles and sceptical arguments: Closure and underdetermination." *Philosophia*, 41, 1125–1133.

BIBLIOGRAPHY

- Brady, Michael and Duncan Pritchard (2005), "Epistemological contextualism: Problems and prospects." *Philosophical Quarterly*, 55, 161–171.
- Brandom, Robert (2008), *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford University Press.
- Brown, Jessica (2005), "Doubt, circularity and the moorean response to the sceptic." *Philosophical Perspectives*, 19, 1–14.
- Brown, Jessica (2013), "Experimental philosophy, contextualism and ssi." *Philosophy and Phenomenological Research*, 86, 233–261.
- Brueckner, Anthony (1994), "The structure of the skeptical argument." *Philosophy and Phenomenological Research*, 54, 827–835.
- Brueckner, Anthony (2005), "Fallibilism, underdetermination, and skepticism." *Philosophy and Phenomenological Research*, 71, 384–391.
- Brueckner, Anthony (2011), " $\sim k \sim sk$." *Philosophical Issues*, 21, 74–89.
- Buchanan, Reid (2002), "Natural doubts: Williams's diagnosis of scepticism." *Synthese*, 131, 57–80.
- Byrne, Alex (2004), "How hard are the sceptical paradoxes?" *Noûs*, 38, 299–325.
- Cappelen, Herman and John Hawthorne (2009), *Relativism and Monadic Truth*. Oxford University Press UK.
- Carruthers, Peter (2009), "How we know our own minds: The relationship between mindreading and metacognition." *Behavioral and Brain Sciences*, 32, 121+.
- Cohen, Stewart (1987), "Knowledge, context, and social standards." *Synthese*, 73, 3–26.
- Cohen, Stewart (1998), "Two kinds of skeptical argument." *Philosophy and Phenomenological Research*, 58, 143–159.
- Cohen, Stewart (1999), "Contextualism, skepticism, and the structure of reasons." *Philosophical Perspectives*, 13, 57–89.
- Cohen, Stewart (2000), "Contextualism and skepticism." *Noûs*, 34, 94–107.
- Cohen, Stewart (2004), "Contextualism and unhappy-face solutions: Reply to schiffer." *Philosophical Studies*, 119, 185–197.

- Craig, Edward (1989), "Nozick and the sceptic: The thumbnail version." *Analysis*, 49, 161–2.
- Craig, Edward (1990), *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.
- Davis, Wayne A. (2004), "Are knowledge claims indexical?" *Erkenntnis*, 61, 257–281.
- Davis, Wayne A. (2007), "Knowledge claims and context: Loose use." *Philosophical Studies*, 132, 395–438.
- DeRose, Keith (1992), "Contextualism and knowledge attributions." *Philosophy and Phenomenological Research*, 52, 913–929.
- DeRose, Keith (1995), "Solving the skeptical problem." *Philosophical Review*, 104, 1–52.
- DeRose, Keith (2004a), "The problem with subject-sensitive invariantism." *Philosophy and Phenomenological Research*, 68, 346–350.
- DeRose, Keith (2004b), "Single scoreboard semantics." *Philosophical Studies*, 119, 1–21.
- DeRose, Keith (2005), "The ordinary language basis for contextualism, and the new invariantism." *Philosophical Quarterly*, 55, 172–198.
- DeRose, Keith (2006), "'bamboozled by our own words': Semantic blindness and some arguments against contextualism." *Philosophy and Phenomenological Research*, 73, 316–338.
- DeRose, Keith (2009), *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*. Oxford University Press.
- Descartes, René (1982), *Discourse on Method and the Meditations*. Penguin Books.
- Dougherty, Trent and Patrick Rysiew (2009), "Fallibilism, epistemic possibility, and concessive knowledge attributions." *Philosophy and Phenomenological Research*, 78, 123–132.
- Dretske, Fred I. (1970), "Epistemic operators." *Journal of Philosophy*, 67, 1007–1023.
- Erana, Angela (2012), "Dual process theories versus massive modularity hypothesis." *Philosophical Psychology*, 25, 855–872.

BIBLIOGRAPHY

- Evans, Jonathan St. B. T. (2007), "On the resolution of conflict in dual process theories of reasoning." *Thinking & Reasoning*, 13, 321–339.
- Evans, Jonathan St B. T. (2009), "How many dual process theories do we need: One, two or many?" In *In Two Minds: Dual Processes and Beyond* (Jonathan Evans and Keith Frankish, eds.), Oxford University Press.
- Evans, Jonathan St B. T. (2011), "Dual-process theories of reasoning: Contemporary issues and developmental applications." *Developmental Review*, 31.
- Evans, Jonathan St B. T. (2014), "Two minds rationality." *Thinking and Reasoning*, 20, 129–146.
- Evans, Jonathan St B. T. and Keith E. Stanovich (2013), "Dual-Process Theories of Higher Cognition: Advancing the Debate." *Perspectives on Psychological Science*, 8, 223–241.
- Faghihi, Usef, Clayton Estey, Ryan McCall, and Stan Franklin (2015), "A cognitive model fleshes out Kahneman's fast and slow systems." *Biologically Inspired Cognitive Architectures*, 11, 38–52.
- Feldman, Richard (1999), "Contextualism and skepticism." *Philosophical Perspectives*, 13, 91–114.
- Feldman, Richard (2001), "Skeptical problems, contextualist solutions." *Philosophical Studies*, 103, 61–85.
- Frankish, Keith (2010), "Dual-process and dual-system theories of reasoning." *Philosophy Compass*, 5, 914–926.
- Frederick, S. (2005), "Cognitive reflection and decision making." *Journal of Economic Perspectives*, 19, 25–42.
- Gerken, Mikkel (2012), "On the cognitive bases of knowledge ascriptions." In *Knowledge Ascriptions* (Jessica Brown and Mikkel Gerken, eds.), Oxford University Press.
- Gerken, Mikkel (2013), "Epistemic focal bias." *Australasian Journal of Philosophy*, 91, 41–61.
- Goldman, A. (2006), *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford University Press.

- Greco, John (2008), "What's wrong with contextualism?" *Philosophical Quarterly*, 58, 416–436.
- Hawthorne, John (2004), *Knowledge and Lotteries*. Oxford University Press.
- Heller, Mark (1999), "Relevant alternatives and closure." *Australasian Journal of Philosophy*, 77, 196–208.
- Herschbach, Mitchell (2015), "Direct social perception and dual process theories of mindreading." *Consciousness and Cognition*, 36, 483–497.
- Hetherington, Stephen Cade (1994), "Sceptical insulation and sceptical objectivity." *Australasian Journal of Philosophy*, 72, 411–425.
- Huemer, Michael (2007), "Epistemic possibility." *Synthese*, 156, 119–142.
- James, William (1911), *Some Problems of Philosophy*. Longmans, Green, and Co.
- Kraft, Tim (2013), "Sceptical scenarios are not error-possibilities." *Erkenntnis*, 78, 59–72.
- Lewis, David (1996), "Elusive knowledge." *Australasian Journal of Philosophy*, 74, 549–567.
- MacFarlane, John (2007), "Relativism and disagreement." *Philosophical Studies*, 132, 17–31.
- Macpherson, Robyn and Keith E. Stanovich (2007), "Cognitive ability, thinking dispositions, and instructional set as predictors of critical thinking." *Learning and Individual Differences*, 17, 115–127.
- McCain, Kevin (2013), "Two skeptical arguments or only one?" *Philosophical Studies*, 164, 289–300.
- McGrath, Matthew (2013), "Dogmatism, underminers and skepticism." *Philosophy and Phenomenological Research*, 86, 533–562.
- Montminy, Martin (2009a), "Contextualism, invariantism and semantic blindness." *Australasian Journal of Philosophy*, 87, 639–657.
- Montminy, Martin (2009b), "Contextualism, relativism and ordinary speakers' judgments." *Philosophical Studies*, 143, 341–356.

- Montminy, Martin (2013), "The role of context in contextualism." *Synthese*, 190, 2341–2366.
- Nagel, Jennifer (2010), "Knowledge ascriptions and the psychological consequences of thinking about error." *Philosophical Quarterly*, 60, 286–306.
- Nagel, Jennifer (2011), "The psychological basis of the harman-vogel paradox." *Philosophers' Imprint*, 11, 1–28.
- Nagel, Jennifer (2012), "Mindreading in gettier cases and skeptical pressure cases." In *Knowledge Ascriptions* (Jessica Brown and Mikkel Gerken, eds.), Oxford University Press.
- Neta, Ram (2003a), "Contextualism and the problem of the external world." *Philosophy and Phenomenological Research*, 66, 1–31.
- Neta, Ram (2003b), "Skepticism, contextualism, and semantic self-knowledge." *Philosophy and Phenomenological Research*, 67, 396–411.
- Nicholls, N. (1999), "Cognitive illusions, heuristics, and climate prediction." *Bulletin of the American Meteorological Society*, 80, 1385–1397.
- Nichols, Shaun and Stephen P. Stich (2003), *Mindreading. An Integrated Account of Pretence, Self-Awareness, and Understanding Other Minds*. Oxford University Press.
- Nozick, Robert (1981), *Philosophical Explanations*. Harvard University Press.
- Oppenheimer, D.M. (2004), "Spontaneous discounting of availability in frequency judgment tasks." *Psychological Science*, 15, 100–105.
- Pritchard, Duncan (2002), "Recent work on radical skepticism." *American Philosophical Quarterly*, 39, 215–257.
- Pritchard, Duncan (2005), *Epistemic Luck*. Clarendon Press.
- Pritchard, Duncan and Chris Ranalli (2013), "Rorty, williams, and davidson: Skepticism and metaepistemology." *Humanities*, 2, 351–368.
- Pryor, James (2000), "The skeptic and the dogmatist." *Noûs*, 34, 517–549.
- Przyjowski, Katrina (2017), "Strong epistemic possibility and evidentiality." *Topoi*, 36, 183–195.

- Putnam, Hilary (1981), *Reason, Truth, and History*. Cambridge University Press.
- Reed, Baron (2013), "Fallibilism, epistemic possibility, and epistemic agency." *Philosophical Issues*, 23, 40–69.
- Rudd, Anthony (2008), "Natural doubts." *Metaphilosophy*, 39, 305–324.
- Rysiew, Patrick (2001), "The context-sensitivity of knowledge attributions." *Noûs*, 35, 477–514.
- Schaffer, Jonathan (2004), "Skepticism, contextualism, and discrimination." *Philosophy and Phenomenological Research*, 69, 138–155.
- Schaffer, Jonathan and Zoltan Gendler Szabo (2013), "Epistemic comparativism: A contextualist semantics for knowledge ascriptions." *Philosophical Studies*, 1–53.
- Schiffer, Stephen (1996), "Contextualist solutions to scepticism." *Proceedings of the Aristotelian Society*, 96, 317–333.
- Schiffer, Stephen (2000), "Vagueness and partial belief." *Noûs*, 34, 220–257.
- Schiffer, Stephen (2004), "Skepticism and the vagaries of justified belief." *Philosophical Studies*, 119, 161–184.
- Sloman, S.A. (1996), "The empirical case for two systems of reasoning." *Psychological Bulletin*, 119, 3–22.
- Smith, E.R. and J. DeCoster (2000), "Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems." *Personality and Social Psychology Review*, 4, 108–131.
- Sosa, Ernest (2009), *Reflective Knowledge*. Oxford University Press.
- Sosa, Ernest and Barry Stroud (1994), "Philosophical scepticism." *Aristotelian Society Supplementary Volume*, 68, 263–307.
- Stanley, Jason (2005), *Knowledge and Practical Interests*. Oxford University Press.
- Stanovich, Keith E. and Maggie E. Toplak (2012), "Defining features versus incidental correlates of type 1 and type 2 processing." *Mind and Society*, 11, 3–13.

BIBLIOGRAPHY

- Stanovich, Keith E., Richard F. West, and Maggie E. Toplak (2011), "The complexity of developmental predictions from dual process models." *Developmental Review*, 31, 103 – 118. Special Issue: Dual-Process Theories of Cognitive Development.
- Stroud, Barry (1984), "Skepticism and the possibility of knowledge." *Journal of Philosophy*, 81, 545–551.
- Stuppel, Edward J. N., Linden J. Ball, and Daniel Ellis (2013), "Matching bias in syllogistic reasoning: Evidence for a dual-process account from response times and confidence ratings." *Thinking & Reasoning*, 19, 54–77.
- Thompson, Valerie A. and Stephen C. Johnson (2014), "Conflict, metacognition, and analytic thinking." *Thinking & Reasoning*, 20, 215–244.
- Thompson, Valerie A., Jamie A. Prowse Turner, and Gordon Pennycook (2011), "Intuition, reason, and metacognition." *Cognitive Psychology*, 63, 107–140.
- Thompson, Valerie A., Jamie A. Prowse Turner, Gordon Pennycook, Linden J. Ball, Hannah Brack, Yael Ophir, and Rakefet Ackerman (2013), "The role of answer fluency and perceptual fluency as metacognitive cues for initiating analytic thinking." *Cognition*, 128, 237–251.
- Toplak, Maggie E., Richard F. West, and Keith E. Stanovich (2014), "Assessing miserly information processing: An expansion of the Cognitive Reflection Test." *Thinking & Reasoning*, 20, 147–168.
- Tversky, Amos and Daniel Kahneman (1973), "Availability: A heuristic for judging frequency and probability." *Cognitive Psychology*, 5, 207–232.
- Tversky, Amos and Daniel Kahneman (1974), "Judgment under uncertainty: Heuristics and biases." *Science*, 185, 1124–1131.
- Tversky, Amos and Daniel Kahneman (1996), "On the reality of cognitive illusions." *Psychological Review*, 103, 582–591.
- Unger, Peter K. (1975), *Ignorance: A Case for Scepticism*. Oxford University Press.
- Vogel, Jonathan (1990), "Are there counterexamples to the closure principle?" In *Doubting: Contemporary Perspectives on Skepticism* (Michael David Roth and Glenn Ross, eds.), 13–29, Dordrecht: Kluwer.

- Willenken, Tim (2011), "Moorean responses to skepticism: A defense." *Philosophical Studies*, 154, 1–25.
- Williams, Michael (1988), "Epistemological realism and the basis of scepticism." *Mind*, 97, 415–439.
- Williams, Michael (1991), *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. B. Blackwell.
- Williams, Michael (2001), "Contextualism, externalism and epistemic standards." *Philosophical Studies*, 103, 1–23.
- Williams, Michael (2012), "Wright against the sceptics." In *Mind, Meaning, and Knowledge: Themes From the Philosophy of Crispin Wright* (Crispin Wright and Annalisa Coliva, eds.), Oxford University Press.
- Williams, Michael (2013), "Skepticism, evidence and entitlement." *Philosophy and Phenomenological Research*, 87, 36–71.
- Williamson, Timothy (2000), *Knowledge and its Limits*. Oxford University Press.
- Williamson, Timothy (2005), "Contextualism, subject-sensitive invariantism and knowledge of knowledge." *The Philosophical Quarterly*, 55, 213–235.
- Wright, Crispin (1991), "Skepticism and dreaming: Imploding the demon." *Mind*, 100, 87–116.
- Wright, Crispin (2002), "(anti-)sceptics simple and subtle: G. e. moore and john mcdowell." *Philosophy and Phenomenological Research*, 65, 330–348.
- Wright, Crispin (2005), "Contextualism and scepticism: Even-handedness, factivity and surreptitiously raising standards." *Philosophical Quarterly*, 55, 236–262.
- Wright, Crispin (2007), "New age relativism and epistemic possibility: The question of evidence." *Philosophical Issues*, 17, 262–283.
- Yarritu, Ion, Helena Matute, and David Luque (2015), "The dark side of cognitive illusions: When an illusory belief interferes with the acquisition of evidence-based knowledge." *British Journal of Psychology*, 106, 597–608.

