Evidentialism, Scepticism and Belief in God

David Austin Hemp

Submitted for the degree of PhD

Department of Philosophy

University of Sheffield

March 2002

Summary

Many people think that it is possible to undermine belief in God by arguing that the evidence for this belief is insufficient. The aim of this thesis is to develop a new objection to this claim (which claim I call *EBG*). In chapter one, I describe an objection to EBG which I call the *anti-sceptical objection*, or the *ASO*. The ASO faces a number of serious problems; but, in later chapters, I show that it can be made invulnerable to these problems.

The ASO can be divided into two stages. Its first stage argues that *EBG is true only if there is a form of evidentialism that discredits belief in God*, and its second stage argues that *there is no form of evidentialism that discredits belief in God.* In chapter two, I argue that the ASO faces a serious problem. This problem is generated by a non-standard form of evidentialism, which I call *explanatory evidentialism.*

In chapter three, I rectify the problem described in chapter two. I do so by constructing a new version of the ASO, which I call *the second ASO.* Unlike the original ASO, the second ASO does not aim to show that EBG is false. Instead, it aims to *undermine* belief in EBG, by showing that, under ordinary standards for knowledge, EBG is not known to be true.

The second ASO can be divided into four stages. Its second stage argues that explanatory evidentialism is the only form of evidentialism which stands a chance of discrediting belief in God. In chapter four, I argue that another form of evidentialism – which I call *epistemic evidentialism* – may discredit belief in God even if explanatory evidentialism fails to do so. Chapter five then constructs a third version of the ASO which is invulnerable to the problem described in chapter four.

Contents

Chapter One

The anti-sceptical objection

Many people think that it is possible to undermine belief in God by arguing that the evidence for this belief is insufficient.[1] The main aim of this thesis is to develop a new objection to this claim. In this chapter, I will outline an objection to this claim which is motivated by reflection on sceptical hypotheses. The objection that I will outline faces a number of serious problems; but, in later chapters, I will show that the objection can be made invulnerable to these problems.

1. Clarifying EBG

Before we can describe the objection with which this chapter is concerned, we must clarify the claim on which this objection is focused (which I will call *EBG*). In what follows, I will make a number of clarifications to the following rough statement of EBG:

> (EBG) It is possible to undermine belief in God by arguing that the evidence for this belief is insufficient.

The clarifications that I will make are motivated by a number of questions and problems that are generated by this rough statement. In the next eight subsections, I will state these questions and problems, and will describe the clarifications that they motivate.

1.1 First clarification

One question that we should ask advocates of (EBG) is the following:

[1] For philosophical defences of this claim, see Clifford 1879, Russell 1957, Blanshard 1974, Flew 1976, and Mackie 1982. For objections to some of these defences, see James 1896, Swinburne 1979, Plantinga 1983 and 2000, and van Inwagen 1997.

(Q1) What do you mean by 'belief in God'?

Different advocates of (EBG) will answer this question in different ways; but, it seems likely that most advocates of the doctrine will want to give an answer which is broad. *Any* belief which says, or entails that there is a deity, or deities will be taken, by most advocates of EBG, to be a belief that can be undermined in the way that EBG describes. For this reason, it is natural to clarify EBG in the following way:

> (EBG_2)It is possible to undermine *any* deity belief by arguing that the evidence for this belief is insufficient.

where a *deity belief* is defined as a belief which says, or entails, that there is a deity, or deities.

1.2 Second clarification

Another question that we could ask an advocate of EBG is the following:

(Q2) What do you mean by 'undermine'?

Most advocates of EBG will want to adopt an epistemic interpretation of this term. On one natural interpretation of this kind, one undermines S's belief that p iff one gives S a *conclusive epistemic reason* to stop believing that p. The best way to clarify the concept of a conclusive epistemic reason is to appeal to examples.

If I make you aware of the fact that your belief that p is false, then I give you a conclusive epistemic reason to stop believing that p. And, if I make you aware of the fact that this belief is epistemically irrational or unreasonable, then I also give you a conclusive epistemic reason to abandon it. There are other ways of giving people conclusive epistemic reasons to abandon their beliefs; but, for the moment, we will not go into the details of these. For now, it is sufficient to say that one gives subject S a conclusive epistemic reason to abandon belief B iff the reason for abandoning B that one gives to S is the same kind of reason as one would give to S by making S aware of the fact that B is epistemically irrational, or false.

On this definition of a conclusive epistemic reason, it seems fair to rewrite (EBG_2) in the following way:

> (EBG_3)It is possible to give those who have deity beliefs a conclusive epistemic reason to abandon these beliefs by arguing that the evidence for these beliefs is insufficient.

For, it seems clear, on this definition, that (EBG_3) states a doctrine that most advocates of EBG would be prepared to endorse.

1.3 Third clarification

A third question that we could ask advocates of EBG is the following:

> (Q3) What, exactly, do you mean by *belief*?

We can motivate this question by focusing on the following example, which is due to Alvin Plantinga:

> 'Consider a Christian beset by doubts. He has a hard time believing certain crucial Christian claims – perhaps the teaching that God was in Christ, reconciling the world to himself. Upon calling that belief to mind, he finds it cold, lifeless, without warmth or attractiveness. Nonetheless, he is committed to this belief; it is his position; if you ask him what he thinks about it, he will unhesitatingly endorse it. He has, so to speak, thrown in his lot with it.'
>
> (Plantinga 1983:37).

Some may be tempted to say that Plantinga's Christian no longer believes the doctrine that Plantinga mentions. But, it seems wrong to make this claim without qualification. For, as Plantinga says, the Christian does still *endorse* this doctrine – it is still his *position*. Because of this, we should say that the doctrine is, in one sense, a doctrine that the Christian believes.

Although the doctrine that Plantinga mentions is, in one sense, a doctrine that the Christian believes, it is not a doctrine in which the Christian has much *confidence*. Upon calling the doctrine to mind, he finds it cold and lifeless. He is by no means sure that it is true. Because of this, we should say that the doctrine is, in one sense, a doctrine that the Christian does *not* believe. By saying this, we commit ourselves to the existence of two different senses of 'believes'.

In what follows, we will say that a subject, S, *actively* believes a proposition, P, iff S believes P in the first of our two senses, and that he *passively* believes it iff he believes it in the second. On this usage, Plantinga's Christian actively believes that God was in Christ, reconciling the world to himself, but does not passively believe that this is so. This use of 'active' and 'passive' is motivated by the fact that, on the first sense of 'believes', belief seems to be a state that can be *actively formed* – i.e. a state that can be formed at will, by performing a certain kind of action. On the second sense of 'believes', active formation of belief does not seem possible, and consequently, it is natural to use the term 'active' for the first sense, and the term 'passive' for the second.[2]

Now that we have distinguished between active and passive belief, we can distinguish between two readings of EBG. On the first reading, EBG says that it is possible to undermine any *active* deity belief by arguing that the evidence for this belief is insufficient, and on the second, it says that it is possible to undermine any *passive* deity belief in this way. On both of these readings, EBG can seem plausible. But, there is reason to think that the first reading gives EBG a better chance of being true than the second.

To see this, let's suppose (a) that I believe, both actively and passively, that *my numbers will not come up in this week's National Lottery draw*, and (b) that my only reason for actively and passively believing this is that *the chance of my numbers*

[2] The distinction between active and passive belief is similar to the distinction that a number of authors have made between belief and acceptance. For some descriptions of this distinction, see de

doctrine that the Christian believes.

Although the doctrine that Plantinga mentions is, in one sense, a doctrine that the Christian believes, it is not a doctrine in which the Christian has much *confidence*. Upon calling the doctrine to mind, he finds it cold and lifeless. He is by no means sure that it is true. Because of this, we should say that the doctrine is, in one sense, a doctrine that the Christian does *not* believe. By saying this, we commit ourselves to the existence of two different senses of 'believes'.

In what follows, we will say that a subject, S, *actively* believes a proposition, P, iff S believes P in the first of our two senses, and that he *passively* believes it iff he believes it in the second. On this usage, Plantinga's Christian actively believes that God was in Christ, reconciling the world to himself, but does not passively believe that this is so. This use of 'active' and 'passive' is motivated by the fact that, on the first sense of 'believes', belief seems to be a state that can be *actively formed* – i.e. a state that can be formed at will, by performing a certain kind of action. On the second sense of 'believes', active formation of belief does not seem possible, and consequently, it is natural to use the term 'active' for the first sense, and the term 'passive' for the second.[2]

Now that we have distinguished between active and passive belief, we can distinguish between two readings of EBG. On the first reading, EBG says that it is possible to undermine any *active* deity belief by arguing that the evidence for this belief is insufficient, and on the second, it says that it is possible to undermine any *passive* deity belief in this way. On both of these readings, EBG can seem plausible. But, there is reason to think that the first reading gives EBG a better chance of being true than the second.

To see this, let's suppose (a) that I believe, both actively and passively, that *my numbers will not come up in this week's National Lottery draw*, and (b) that my only reason for actively and passively believing this is that *the chance of my numbers*

[2] The distinction between active and passive belief is similar to the distinction that a number of authors have made between belief and acceptance. For some descriptions of this distinction, see de

coming up is extremely small. In these circumstances, you can give me a conclusive epistemic reason to abandon my active belief about the lottery by arguing that my evidence for this belief is insufficient. But, it is not clear that you can give me a conclusive epistemic reason to abandon my passive belief about the lottery in this way. To see this, recall that passive belief is a matter of confidence. To passively believe a proposition is, roughly, to be very confident that that proposition is true. It seems clear that it is reasonable, in this situation, for me to be very confident that my numbers will not come up. And, consequently, it seems clear that you cannot give me a conclusive epistemic reason to stop passively believing that my numbers will not come up by arguing that the evidence for this belief is insufficient.

The overall aim of this thesis is to develop and defend a certain objection to EBG. Because of this, we should interpret EBG as charitably as possible. The above comments suggest that the most charitable way of interpreting EBG is as a claim about active, rather than passive belief. So, in what follows, we will interpret EBG in this way.

To save words, I will refrain from explicitly restating EBG as a claim about active deity beliefs. But, in what follows, it should be read as a claim of this kind. More generally: *any* claim that I make about belief *simpliciter* should be read as a claim about active belief. Claims that I make about passive belief will be explicitly stated as such.

1.4 Fourth clarification

Many people who have deity beliefs are aware of the apparent lack of evidence for these beliefs. Advocates of EBG may want to say that some or all of these people are already aware of the fact that the evidence for their deity beliefs is insufficient. If a subject, S, is already aware of the fact that the evidence for his deity beliefs is insufficient, then it will not be possible to *give* S a reason to abandon his deity beliefs by arguing for this fact. Because of this, section 1.2's statement of EBG needs to be amended, in something like the following way:

Sousa 1971, Cohen 1992 and Frankish 1998.

(EBG_4)If there are people who have deity beliefs, and who are not aware of the fact that the evidence for these beliefs is insufficient, then we can give these people a conclusive epistemic reason to abandon their deity beliefs, by arguing that the evidence for these beliefs is insufficient.

1.5 Fifth clarification

Some advocates of EBG may not want to commit themselves to saying that the evidence for people's deity beliefs is *in fact* insufficient. They may want to restrict themselves to endorsing the weaker claim that this evidence is not *known* by us to be sufficient, or that is cannot *reasonably be believed* by us to be so. To accommodate these people, we should amend (EBG_4) in something like the following way:

(EBG_5)If there are people who have deity beliefs, and who are not aware of the fact that the following claim (or some suitable variant of this claim) is true with respect to them:

(¬E) The evidence for my deity beliefs is insufficient.

then we can give these people a conclusive epistemic reason to abandon their deity beliefs by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them.

Once we have amended it in this way, we can say that claims like the following are suitable variants of (¬E):

(¬KE) I do not know that the evidence for my deity beliefs is sufficient.

(¬RBE) It is not reasonable for me to believe that the evidence for my deity beliefs is sufficient.

We can also classify claims like the following as suitable variants of (¬E):

(¬HE)	I do not have sufficient evidence for my deity beliefs.
(¬BE)	My deity beliefs are not based on sufficient evidence.
(¬KBE)	I do not know whether my deity beliefs are based on sufficient evidence.

However: we should not classify claims like the following as suitable variants of (¬E):

(EA)	There is evidence against my deity beliefs.
(EAF)	The evidence against my deity beliefs outweighs the evidence for them.

For, the core idea of EBG is that we can undermine deity beliefs by arguing that there is insufficient evidence *for* them. Advocates of EBG may believe that there is evidence against the existence of deities, and they may believe that this evidence can be used to undermine deity beliefs. But, they do not take themselves to be appealing to evidence of this kind when they endorse EBG; so, evidence of this kind should not be mentioned in our statement of EBG.

1.6 Sixth clarification

A belief, B, is a deity belief iff B says, or entails, that there is a deity, or deities. It is useful to distinguish between two kinds of deity belief. On the one hand, there are *epistemic* deity beliefs, which say, or entail that we *know* something about a deity, or deities. And, on the other, there are *metaphysical* deity beliefs, which do not say or entail anything of this kind.

Many advocates of EBG hold that all metaphysical deity beliefs (henceforth: MDBs) are untrue. But, no-one endorses EBG *because* they believe that all MDBs are untrue. Those who endorse EBG would endorse it even if they were agnostic about the truth-value of MDBs. Because of this, they should be prepared to endorse the following strengthened formulation of EBG:

(EBG_6) If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them then we can give these people a conclusive epistemic reason to abandon their deity beliefs by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them. *We can this regardless of whether the MDBs of these people are untrue.*

It is important that the italicised clause of this formulation refers only to *metaphysical* deity beliefs. For, advocates of EBG would *not* endorse this doctrine if they were agnostic about the truth-value of *epistemic* deity beliefs. To see this, we need to think about the reasons that advocates of EBG have for endorsing this doctrine. It seems clear that those who endorse (EBG_6) do so because they believe something like the following claim:

(EBG_6') If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them, then we can *make these people aware of the fact that their deity beliefs have some epistemic defect, D,* by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them. We can do this regardless of whether the MDBs of these people are untrue.

There may be some disagreement among advocates of EBG as to the nature of epistemic defect D. Some may take it to be the defect of irrationality, while others take it to be some less serious defect. But, all of them are likely to agree that beliefs which have defect D do not constitute knowledge. So, it is likely that they will all endorse the following claim:

(EBG_6") If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them then we can make these people aware of the fact that their deity beliefs *do not constitute knowledge* by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them. We can do this regardless of whether the MDBs of these people are untrue.

Anyone who endorses this claim will think that there are no deity beliefs that constitute knowledge. Consequently, they will think that no epistemic deity beliefs (henceforth: EDBs) are true. From this, it follows that, if advocates of EBG were agnostic about the truth-value of some EDBs, then they would not endorse (EBG_6"). And, if advocates of EBG did not endorse (EBG_6"), then they would not endorse (EBG_6); so, we can conclude that advocates of EBG would not endorse EBG if they were agnostic about the truth-value of some epistemic deity beliefs.

1.7 Seventh clarification

Because advocates of EBG endorse (EBG_6"), they are committed to endorsing the following claim:

(EBG_{6K}) If there are people who have deity beliefs, and who believe that these deity beliefs constitute knowledge, then we can make these people aware of the fact that their deity beliefs do not constitute knowledge, by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them.

Some advocates of EBG may think that there are *other* ways of making people aware of the fact that their deity beliefs do not constitute knowledge. E.g. they may hold that we can make people aware of this fact by arguing that their deity beliefs were *caused in the wrong way*, or by arguing that these beliefs were produced by a *cognitive malfunction*. If they do hold this, then they hold that certain epistemic deity beliefs can be rebutted in a *non-evidential* way.

Even if some advocates of EBG do believe that certain EDBs can be rebutted in a non-evidential way, it seems clear that they do not endorse EBG *because* they believe this to be so. Those who endorse EBG would endorse it even if they were agnostic about whether there is a non-evidential way of rebutting EDBs. Because of this, they should be prepared to endorse the following, strengthened formulation of EBG:

(EBG_7) If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them, then we can give these people a conclusive epistemic reason to abandon their deity beliefs by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them. *We can do this regardless of whether the MDBs of these people are untrue, and regardless of whether their EDBs can be rebutted in a non-evidential way.*

The objection to EBG that we will state in this chapter trades on the fact that advocates of EBG are committed to the italicised clause of the above formulation.

1.8 Eighth clarification

Some people think that we can undermine just about any belief by arguing that the evidence for that belief is insufficient. These people endorse a strong form of scepticism. Others think that, although scepticism of this kind is not ordinarily true, there are certain unusual contexts in which the standards for epistemic appraisal are raised in such a way as to make it true.[3] These people endorse a form of epistemic contextualism.

Most advocates of EBG do not endorse it for sceptical, or contextualist reasons. They do not endorse EBG because they think that similar claims hold for just about any belief, and they do not endorse it because they think that we can raise the standards of epistemic appraisal in such a way as to make EBG true. According to them, EBG holds under ordinary standards of epistemic appraisal, and holds even if global scepticism is false. Because of this, they should be prepared to endorse the following strengthened formulation of EBG:

> (EBG_8)If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them, then we can give these people a conclusive epistemic reason to abandon their deity beliefs by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them. We can do this regardless of whether the MDBs of these people are untrue, and regardless of whether their EDBs can be rebutted in a non-evidential way. *We can also do it regardless of whether global scepticism is true, and regardless of whether the standards for epistemic appraisal have been raised.*

In what follows, references to EBG should be interpreted as references to (EBG_8).

2. The anti-sceptical objection

Now that we have clarified EBG, we can start to describe the objection to EBG with which this chapter is concerned.[4] Because this objection is motivated by reflection on certain sceptical hypotheses, we will call it *the anti-sceptical objection*, or *ASO*. The first stage of the ASO is an argument for the following claim:

> (1) EBG is true only if there is a form of evidentialism that discredits all deity beliefs.

And, the second stage is an argument for this claim:

> (2) There is no form of evidentialism that discredits all deity beliefs.

Before we can start to state these arguments, we need to explain what is meant by 'a

[3] See e.g. Cohen 1988 and 1999.
[4] The objection that we will describe is inspired by the arguments in Plantinga 1983 and Plantinga

form of evidentialism', and what it is for a form of evidentialism to *discredit* all deity beliefs. Section 3 will explain what is meant by these terms, and sections 4 and 5 will give the arguments for claims (1) and (2) from which the anti-sceptical objection is composed.

3. Forms of evidentialism

By a form of evidentialism, I mean a doctrine that is sufficiently similar to the following:[5]

> (E_1) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P is epistemically irrational.

A doctrine is sufficiently similar to (E_1) iff it is either identical to (E_1), or can be generated by making an *admissible amendment* to (E_1). There are four kinds of admissible amendment that can be made to (E_1). In the next four subsections, I will say something about each kind of amendment, and in section 3.5, I will explain what it is for a form of evidentialism to discredit all deity beliefs.

3.1 Amending the doxastic defect

One way of amending (E_1) is to amend the doxastic defect to which it refers. The following doctrines can both be generated by an amendment of this kind:

> (E_2) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P is *false*.

> (E_3) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then

1985. It also owes something to the arguments in van Inwagen 1997 and DeRose 2000.

[5] My use of this term is similar to the use adopted in Feldman and Conee 1985.

S's belief in P *does not constitute knowledge*.

Amendments of this kind are admissible iff the defect to which they appeal is *suitable*, where a doxastic defect, D, is said to be suitable iff it satisfies the following constraint:

(SDD) For any subject, S, and proposition P: if S believes P, and S's belief in P has defect D, then the fact that this belief has D is a conclusive epistemic reason for S to stop believing P.

The amendments which generate (E_2) and (E_3) both seem to be admissible; but it is not clear that there are any other admissible amendments of this kind. For, it seems clear that a doxastic defect will be suitable iff it is epistemic. And, it is not clear that there are any epistemic defects that differ significantly from the defects to which (E_1) – (E_3) appeal.

3.2 Amending the evidential requirement

Another way of amending (E_1) is to make changes to the evidential requirement that it imposes. We can generate the following doctrines by making amendments of this kind:

(E_4) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and *there is* not enough evidence for P, then S's belief in P is epistemically irrational.

(E_5) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not *based* on sufficient evidence, then S's belief in P is epistemically irrational.

(E_6) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S *does not know* that his belief in P is based on sufficient evidence, then S's belief in P is epistemically irrational.

Amendments of this kind are admissible iff they do not depart from the key idea that many of our beliefs require *the presence of positive evidence*, in order to be free from epistemic defects. The amendments that generate (E_4) – (E_6) all adhere to this key idea. But, the amendments that generate the following doctrines do not:

(E_7) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and there is evidence *against* P, then S's belief in P is epistemically irrational.

(E_8) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and *S has* evidence against P, then S's belief in P is epistemically irrational.

(E_9) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S *knows* that she has evidence against P, then S's belief in P is epistemically irrational.

These doctrines are all motivated by the idea that many beliefs require the *absence of negative* evidence, in order to be free from epistemic defect. And, this idea is very different from the evidentialist idea encapsulated in doctrines like (E_4) –(E_6).

3.3 Substituting alleged synonyms

A third way of amending (E_1) is to replace its constituent terms with terms that are, at least allegedly, its synonyms. All of the following doctrines can be generated by making amendments of this kind:

(E_{10}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have *rational perceptual beliefs that confirm P to a sufficient degree*, then S's belief in P is epistemically irrational.

(E_{11}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have *some* rational beliefs that confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

(E_{12}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does have any *knowledge-constituting beliefs* that confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

The above amendments are all generated by replacing 'has sufficient evidence for P' with alleged synonyms; but, similar amendments could be generated by doing the same with respect to 'is epistemically irrational', or with respect to other terms in (E_1). Amendments of this kind are admissible iff they replace terms in (E_1) with terms that can reasonably be believed to be their synonyms. There are many admissible amendments of this kind.

3.4 Making complex amendments

A fourth way of amending (E_1) is to make an admissible amendment to a doctrine that has itself been generated by making an admissible amendment to (E_1). The following doctrines can all be generated by making amendments of this kind:

(E_{13}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not *based on* rational perceptual beliefs that confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

(E_{14}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not based on knowledge-constituting beliefs that confirm P to a sufficient degree, then S's belief in P does not constitute knowledge.

(E_{15}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not ultimately based on perceptual beliefs that confirm P to a sufficient degree, then S's belief in P does not constitute knowledge.

Every amendment of this kind is an admissible amendment to (E_1). As before, there are many admissible amendments of this kind.

3.5 Discrediting theistic belief

The last four subsections explain what it is for an amendment to (E_1) to be admissible. By doing so, they show us what it is for a doctrine to be a form of evidentialism. In this section, I will explain what it is for a form of evidentialism to discredit all deity beliefs. The next section will argue that EBG is true only if there is a form of evidentialism that discredits all deity beliefs.

A form of evidentialism, E, discredits all deity beliefs iff it is the case that, for every subject, S, who has deity beliefs, the following three claims are all true:

(D1) S is in a position to know E, and is in a position to know that her deity beliefs are not exceptions to it.

(D2) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that E imposes, then we can make S aware of this fact.

(D3) (D1) and (D2) are true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way. They are also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

On this definition, it is clear that the claim that *there is a form of evidentialism that discredits all deity beliefs* is very similar to the claim that we are calling EBG. In section 4, I will use the similarities between these two claims to argue for the following claim:

(1) EBG is true only if there is a form of evidentialism that discredits all deity beliefs.

The argument of section 4 is the first stage of the anti-sceptical objection.

4. Arguing for (1)

In section 4.1, I will argue that EBG implies the following claim:

(EBG') If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them, then we can *make these people aware of the fact that their deity beliefs have a suitable doxastic defect* by arguing that (¬E) (or some suitable variant of (¬E)) is true with respect to them. We can do this regardless of whether the MDBs of these people are untrue, and regardless of whether their EDBs can be rebutted in a non-evidential way. We can also do it regardless of whether global scepticism is true, and regardless of whether the standards for epistemic appraisal have been raised.

Section 4.2 will argue that (EBG') itself implies the following claim:

(EBG") If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them, then (i) there is an *evidentialist conditional*, E_C, that these people are in a position to know, (ii) we can make these people aware of the truth of the antecedent of E_C, (iii) claims (i) and (ii) are true regardless of whether the MDBs of these people are untrue, and regardless of whether their EDBs can be rebutted in a non-evidential way, and (iv) claims (i) and (ii) are true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic appraisal have been raised.

And, section 4.3 will argue that, if (EBG'') is true, then there is a form of evidentialism that discredits all deity beliefs. If the arguments of sections 4.1 to 4.3 succeed, then they will show that *EBG is true only if there is a form of evidentialism that discredits all deity beliefs*. In section 5, we will describe an argument which aims to show that *there is no form of evidentialism that discredits all deity beliefs*.

4.1 From EBG to (EBG')

The best way of arguing that EBG implies (EBG') is to appeal to the following claim:

> (C) The only way of giving a subject, S, a conclusive epistemic reason to abandon a belief, B, is to make S aware of the fact that B has some kind of *defect*.

This claim is very plausible. And, it is easy to argue that, if this claim is true, then EBG implies (EBG'). The only thing that we need to do, in order to argue for this implication, is to point out that (C) entails the following claim (where suitability of doxastic defects is defined in the same way as it was in section 3.1):

> (C) The only way of giving a subject, S, a conclusive epistemic reason to abandon a belief, B, is to make S aware of the fact that B has a *suitable* defect.

Given our definition of suitability, it is clear that the defect to which (C) refers must be suitable. And, once we see that this defect must be suitable, it is easy to see that, if (C) is true, then EBG implies (EBG').

4.2 From (EBG') to (EBG'')

Before we can argue that (EBG') implies (EBG''), we need to introduce the concept of an *evidentialist conditional*. By an evidentialist conditional, I mean a conditional whose antecedent is the following claim (or is some suitable variant of this claim):

(¬E) The evidence for my deity beliefs is insufficient.

and whose consequent is a claim of the following form (where defect D is a suitable doxastic defect):

(DBD) My deity beliefs have defect D.

If the term 'evidentialist conditional' is defined in this way, then (EBG') is equivalent to the following claim:

(EBG'*) If there are people who have deity beliefs, and who are not aware of the fact that (¬E) (or some suitable variant of (¬E)) is true with respect to them, then we can make these people aware of the fact that the consequent of a certain evidentialist conditional, E_C, is true with respect to them, by arguing that the antecedent of E_C is true with respect to them. We can do this regardless of whether the MDBs of these people are untrue, and regardless of whether their EDBs can be rebutted in a non-evidential way. We can also do it regardless of whether global scepticism is true, and regardless of whether the standards for epistemic appraisal have been raised.

And, if (EBG') is equivalent to this claim, then it surely implies (EBG''). For, it seems generally true that, if we can make some subject, S, aware of the truth of conditional C's consequent by arguing for conditional C's antecedent, then S is in a position to know C. And, if this is generally true, then (EBG'*) clearly implies (EBG'').

4.3 (EBG″) and evidential discredit

If (EBG″) is true, then there is a form of evidentialism that discredits all deity beliefs. To see this, we need to think about the relationship between evidentialist conditionals and forms of evidentialism. It seems clear that the only way in which a subject, S, can come to know an evidentialist conditional is by inferring it from some form of evidentialism that is a conceptual truth. And, it seem equally clear that, if there is some form of evidentialism, E, that is a conceptual truth, then we are all in a position to know E. From these two points, it follows that, if (EBG″) is true, then, there is some form of evidentialism, E, which is such that, for every subject, S, who has deity beliefs, the following three claims are true:

(D1) S is in a position to know E, and is in a position to know that her deity beliefs are not exceptions to it.

(D2) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that E imposes, then we can make S aware of this fact.

(D3) (D1) and (D2) are true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way. They are also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic appraisal have been raised.

And, from this fact, and the definitions of section 3.5, we can infer that, if (EBG″) is true, then there is a form of evidentialism that discredits all deity beliefs.

5. Arguing for (2)

In the last section, we stated the first stage of the anti-sceptical objection, which is an argument for the following claim:

(1) EBG is true only if there is a form of evidentialism that discredits all deity beliefs.

In this section, we will state the second stage of this objection, which is an argument for this claim:

(2) There is no form of evidentialism that discredits all deity beliefs.

One good way of stating this argument is to focus on the following form of evidentialism, which we introduced at the start of section 3:

(E_1) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P is epistemically irrational.

In section 5.1, we will argue that (E_1) does not discredit all deity beliefs. Sections 5.2 – 5.5 will then argue that there is no admissible amendment to (E_1) which makes it more likely to discredit all deity beliefs. If there is no amendment of this kind, then there is no form of evidentialism that discredits all deity beliefs. So, if the arguments of sections 5.1 - 5.5 all succeed, then they will show that claim (2), above, is true.

5.1 Against (E_1)

(E_1) cannot discredit all deity beliefs unless (E_1) is true.[6] In this section, I will give an argument which aims to show that (E_1) is not true.[7] The gist of this argument is that (E_1) leads to scepticism. One good way of stating the argument is to focus on the sceptical hypothesis that I am a brain in a vat (henceforth: the hypothesis that *I am a BIV*)[8].

[6] To see this, recall that a form of evidentialism, E, discredits all deity beliefs only if those who have deity beliefs are in a position to *know* E.

[7] This argument is inspired by a similar argument in DeRose 2000.

[8] By the hypothesis that *I am a brain in a vat*, I mean the hypothesis that I am a disembodied brain floating in a vat of nutrient fluid and being electrochemically stimulated to have exactly the experiences that I am actually having.

The negation of the BIV hypothesis is not a self-evident truth, like the proposition that everything is self-identical. For this reason, it is plausible to suppose that the negation of this hypothesis is not one of the exceptional propositions to which the qualification at the start of (E_1) refers. If the negation of the BIV hypothesis is not one of these exceptional propositions, then (E_1) commits us to endorsing the following claim:

> (E_{1BIV}) If I do not have enough evidence for the proposition that *I am not a BIV*, then my belief that *I am not a BIV* is epistemically irrational.

There is good reason to think that this claim commits us to an implausible form of scepticism.

To see this, note first of all that, if the proposition that *I am not a BIV* was false (i.e. if I was a BIV), then things would seem to me exactly as they actually do. The fact that things would seem this way seems to show that I do not have *any* evidence for the proposition that *I am not a BIV.* But, if I do not have any evidence for this proposition, then (E_{1BIV}) commits us to saying that my belief in this proposition is epistemically irrational. And, if my belief that *I am not a BIV* is epistemically irrational, then the same goes for almost all of my other beliefs.

To see this, consider my belief that *I have hands*. The proposition that *I have hands* obviously entails that *I am not a BIV* – so, if it is epistemically rational for me to believe the first proposition, then it is surely rational for me to believe the second. From this, it follows that, if it is epistemically *irrational* for me to believe that I am not a BIV, then it is similarly irrational for me to believe that I have hands. Reasoning of this kind will establish similar conclusions with respect to almost all of my perceptual beliefs; so, if we endorse (E_1), then we seem committed to endorsing a very strong form of scepticism.

Advocates of (E_1) might try to avoid this problem by insisting that, contrary to appearances, the negation of the BIV hypothesis *is* one of the exceptional propositions to which the qualification in (E_1) refers. But, this strategy does not seem promising. For, we can restate the argument of the last two paragraphs by appealing to other sceptical hypotheses. One way in which we could do this is by appealing to other 'global' sceptical hypotheses, like the following:

(IND) My expectations about the future will all be confounded by a breakdown in the law-like regularities on which these expectations are all based.

(MIND) All of the 'people' whom I know are in fact cleverly designed automata, which act just as if they have minds, but which are in fact completely mindless.

(MEM) The world came into existence five minutes ago, complete with all of the apparent traces of the past on which my beliefs about the past are based.

(TEST) Everything that I believe on the basis of testimony is a lie, which has been told to me as part of an elaborate conspiracy that is designed to prevent me from discovering the true nature of the world.

Another equally effective strategy is to appeal to 'local' sceptical hypotheses, like the following:

(WALL) The visual experiences on the basis of which I believe that *the library walls are blue* are not veridical. In fact, the library walls are white, but the lighting within the library has been cleverly arranged in such a way as to make them look blue.

(THERM) The thermometer reading on the basis of which I believe that *my temperature is normal* was inaccurate. In fact, I have a high temperature, but my thermometer is not working properly, and so it failed to register this.

(NEWS) The newspaper listings on the basis of which I believe that *Spurs beat Fulham* were inaccurate. In fact, Fulham beat Spurs, but a misprint in my newspaper has led me to believe that the opposite is true.

With regard to each of these hypotheses, it is plausible to say that, if the hypothesis was true, then things would – or at least, might – seem to me just as they actually do. Consequently, it is plausible to say that I do not have evidence – or at least, not *sufficient* evidence – for the falsity of any of these hypotheses. It is also plausible to say that the negations of these hypotheses are not exceptions to the universal generalisation that appears in (E_1). And so, we can use each of these hypotheses to argue that (E_1) commits us to an implausible form of scepticism.[9]

The above comments strongly suggest that there are many sceptical hypotheses that can be used to attack (E_1). If there are many hypotheses of this kind, then we cannot save (E_1) by insisting that the negations of these hypotheses are all among the exceptional propositions to which the qualification in (E_1) refers. For, according to this qualification, there are no more than a few exceptions to the generalisation in (E_1). Because of this, we cannot say that every sceptical hypothesis is an exception to the generalisation in (E_1) without conceding that (E_1) is false.

If the argument of this subsection succeeds, then it shows that (E_1) is false. And, if it shows that (E_1) is false, then it shows that (E_1) does not discredit all deity beliefs. In the next five subsections, I will argue that there are no admissible amendments to (E_1) which make it more likely to discredit all deity beliefs. If the arguments of these subsections succeed, then they will show that there is *no* form of evidentialism that discredits all deity beliefs.

5.2 Amending the doxastic defect

One way of amending (E_1) is to amend the doxastic defect to which it refers. Such amendments are admissible iff they replace this defect with a defect that is suitable – where a doxastic defect, D, is suitable iff it satisfies the following constraint:

---

[9] The forms of scepticism to which the local sceptical hypotheses commit us will not be as far-reaching as those which the global hypotheses seem to support. But, they will nevertheless be implausible. It is clearly rational for me to believe that *the walls in the library are blue*, even if I have not checked for the kind of tricks which the hypothesis (WALL) describes. (E_1) seems to entail that it

(SDD) For any subject, S, and proposition P: if S believes P, and S's belief in P has defect D, then the fact that this belief has D is a conclusive epistemic reason for S to stop believing P.

On this definition of suitability, there seem to be just three defects that are suitable. The first is the defect of *falsity*, the second is the defect of *epistemic irrationality*, and the third is the defect of *not constituting knowledge*. Consequently, there seem to be just two admissible ways of amending the defect to which (E_1) refers. The first is to rewrite (E_1) in the following way:

(E_2) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P is *false*.

and the second is to rewrite it as follows:

(E_3) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P *does not constitute knowledge*.

Neither of these amendments turns (E_1) into a doctrine that discredits all deity beliefs. For, neither turns it into a doctrine that is true. (E_2) is obviously untrue; for, there are clearly many propositions that can be believed truly without evidence. And, although (E_3) is not obviously untrue, it can be shown to be untrue by an amended version of the argument of section 5.1. To see this, we need only note that the argument of section 5.1 says very little about epistemic rationality. The only assumptions about epistemic rationality that advocates of this argument need to make are (i) that epistemic rationality is preserved by deductive inference, and (ii) that many of our beliefs about the external world are epistemically rational. Both of these assumptions seem just as plausible when restated in terms of knowledge. So, by restating the argument of section 5.1, we can show that, like (E_1), (E_3) is false.

---

is *not* rational for me to believe this; and consequently, we seem to have reason to reject (E_1).

The above considerations seem to show that we cannot generate a form of evidentialism that discredits all deity beliefs by amending the doxastic defect to which (E_1) refers. In section 5.3, I will argue that amendments to the evidential requirement in (E_1) are no more likely to generate a doctrine that discredits all deity beliefs.

5.3 Amending the evidential requirement

There are two kinds of admissible amendment that can be made to the evidential requirement that is imposed by (E_1). Amendments of the first kind strengthen this requirement, by rewriting (E_1) in something like the following way:

> (E_4) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not *based* on sufficient evidence, then S's belief in P is epistemically irrational.

> (E_5) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S *does not know* that her belief in P is based on sufficient evidence, then S's belief in P is epistemically irrational.

And, amendments of the second kind weaken it, by rewriting (E_1) in something like this way:

> (E_6) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have *any* evidence for P, then S's belief in P is epistemically irrational.

> (E_7) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and *there is* not enough evidence for P, then S's belief in P is epistemically irrational.

Amendments of the first kind cannot turn (E_1) into a doctrine that discredits all deity beliefs. For, every form of evidentialism that is generated by an amendment of this kind will entail (E_1), and will thus be vulnerable to the argument of section 5.1. The only way of saving (E_1) from this argument is to weaken the evidential requirement that it imposes on our beliefs. But, there is reason to think that, even when (E_1) is amended in this way, it will not discredit all deity beliefs.

To see this, note first of all that, if we want to make (E_1) invulnerable to the argument of section 5.1, then we will have to weaken it in something like the following way[10]:

(E_7) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and *there is* not enough evidence for P, then S's belief in P is epistemically irrational.

If (E_1) is amended in this way, then it may not be vulnerable to the argument of section 5.1. For, while it is plausible to suppose that *I do not have* any evidence against the sceptical hypotheses that are mentioned in this section, it is by no means clear that *there is* no evidence against these hypotheses. If the hypotheses are in fact false, then there will be plenty of evidence against them: the fact that *I have hands* will be evidence against the BIV hypothesis; the fact that *the sun will rise tomorrow* will be evidence against the hypothesis that we called (IND), and so on. So, if we rewrite (E_1) as (E_7), then we may make it invulnerable to the argument of section 5.1.

However: even if (E_7) is invulnerable to the argument of section 5.1, there is good reason to believe that (E_7) does not discredit all deity beliefs. To see this, recall that (E_7) discredits all deity beliefs only if it is the case that, for every subject, S, who has deity beliefs, the following claims are *both* true:

($D2_{E7}$) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that (E_7) imposes, then we can make S aware of this fact.

($D3_{E7}$) ($D2_{E7}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way.

The first of these claims may be true with regard to every subject, S, who has deity beliefs. But, there is good reason to believe that the second is not. One good way of seeing this is to note that many of those who have deity beliefs believe the doctrine of *theism*. By the doctrine of theism, I mean the doctrine that there is an omnipotent, omniscient, morally perfect God who created the universe, and who keeps it in existence.

If the doctrine of theism is true, then there is plenty of evidence for its truth. Among this evidence is the fact that *God believes that theism is true*, which, given God's omniscience, provides excellent evidence for the truth of theism. Because the truth of theism implies that there is plenty of evidence for theism, any subject, S, who believes the doctrine of theism is such that, if S's metaphysical deity beliefs are all true, then ($D2_{E7}$) is not true with respect to S. From this, it follows that, for every subject, S, who believes the doctrine of theism, ($D3_{E7}$) is not true.

The above considerations seem to show that we cannot generate a form of evidentialism that discredits all deity beliefs by amending the evidential requirement that (E_1) imposes. In section 5.4, I will argue that the substitution of terms in (E_1) for terms that are allegedly their synonyms is no more likely to generate a form of evidentialism that discredits all deity beliefs.

5.4 *Substitution of alleged synonyms*

There are many kinds of admissible amendment to (E_1) that can be made by substitution of alleged synonyms. Some of these amendments seem to generate forms of evidentialism that are invulnerable to the argument of section 5.1. In what follows, I will consider a representative sample of the forms of evidentialism that can be generated by amendments of this kind. By focusing on the members of this sample, I will argue that no form of evidentialism that is generated by an amendment

[10] To see this, note that (E_6) is as vulnerable as (E_1) to the argument of section 5.1.

of this kind is capable of discrediting all deity beliefs.

One form of evidentialism that can be generated by substitution of alleged synonyms is the following:

(E_8) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and the propositions that S knows to be true do not confirm P to a sufficient degree, then S's belief in P is epistemically irrational.[11]

This form of evidentialism seems to be invulnerable to the argument of section 5.1. For, all of us seem to know propositions which *entail* – and thus strongly confirm - the negations of the sceptical hypotheses on which this argument focuses.[12]

Although (E_8) seems invulnerable to the argument of section 5.1, there is good reason to believe that (E_8) does not discredit all deity beliefs. To see this, recall that (E_8) discredits all deity beliefs only if it is the case that, for every subject, S, who has deity beliefs, the following pair of claims are true:

($D2_{E8}$) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that (E_8) imposes, then we can make S aware of this fact.

($D3_{E8}$) ($D2_{E8}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way.

It seems clear that, to make a subject, S, aware of the fact that her deity beliefs don't satisfy the evidential requirement that (E_8) imposes, we must first make S aware of the fact that her deity beliefs do not constitute knowledge. But, if we can make S aware of the fact that her deity beliefs do not constitute knowledge *before* we have

[11] If the arguments of Williamson 1997 succeed, then they show that (E_1) is synonymous with (E_8).
[12] To see this, note (a) that I seem to know *that I have hands*, which proposition entails that I am not a BIV; (b) that I seem to know *that the sun will rise tomorrow*, which proposition entails the negation of (IND); and likewise for each of the hypotheses that were described in section 5.1.

made aware of the fact that her deity beliefs don't satisfy the evidential requirement that (E_8) imposes, then we can rebut some of S's epistemic deity beliefs in a non-evidential way. Because of this, it seems fair to conclude that there are at least some subject for whom ($D2_{E8}$) and ($D3_{E8}$) are not both true.

If there are some subjects for whom these claims are not both true, then (E_8) does not discredit all deity beliefs. Are there any other forms of evidentialism which can be generated by the substitution of alleged synonyms, and which stand a better chance of discrediting all deity beliefs? One form of evidentialism which might be thought to stand a better chance of discrediting such beliefs is the following:

(E_9) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and the propositions that S knows *on the basis of perception* do not confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

Like (E_8), this form of evidentialism can be generated by substitution of alleged synonyms. And, like (E_8), it seems to avoid some of the sceptical consequences of (E_1). (E_1) commits us to saying that my belief that *I am not a BIV* is epistemically irrational. But, (E_9) does not seem to commit us to saying this: for, I seem to have perceptual knowledge of many propositions that entail that *I am not a BIV*.

If (E_9) avoids all of the sceptical consequences of (E_1), then it may discredit all deity beliefs. For, deity beliefs are not generally confirmed to any great degree by perceptual beliefs; so, it may turn out that, for every subject, S, who has deity beliefs, the following pair of claims are true:

($D2_{E9}$) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that (E_9) imposes, then we can make S aware of this fact.

($D3_{E9}$) ($D2_{E9}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way.

But, there is reason to think that (E_9) does not avoid all of the sceptical consequences of (E_1). One good way of seeing this is to focus on one of the other sceptical hypotheses that we described in section 5.1. The negation of the following sceptical hypothesis is entailed by each of the propositions that I believe on the basis of testimony:

(TEST) Everything that I believe on the basis of testimony is a lie, which has been told to me as part of an elaborate conspiracy that is designed to prevent me from discovering the true nature of the world.

So, if my belief in the negation of (TEST) is epistemically irrational, then the same goes for all of the beliefs that I hold on the basis of testimony. The propositions that I know on the basis of perception seem to be propositions that would have been true even if (TEST) had been true; and, this strongly suggests that these propositions do not confirm the negation of (TEST) to a sufficient degree. Consequently, there is reason to think that, if (E_9) is true, then all of the beliefs that I hold on the basis of testimony are epistemically irrational.

By appealing to our discussion of (E_9), we can argue that no evidentialist doctrine which has the following form will discredit all deity beliefs:

(E_x) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and the propositions that S knows on the basis of X do not confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

We can argue for this by arguing that every evidentialist doctrine, E_x, which has this form either leads to scepticism, or is such that, for some subject, S, the following pair of claims are not both true:

($D2_{EX}$) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that E_x imposes, then we can make S aware of this fact.

($D3_{EX}$) ($D2_{EX}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way.

The instance of (E_x) which replaces 'X' with 'perception' leads to scepticism, as we have just seen. The instance that replaces 'X' with 'perception and testimony' may not lead to scepticism; but, since there are many people who hold some deity beliefs on the basis of testimony, it seems likely that this instance of (E_x) – which we can call (E_{10}) – will be such that, for some subject, S, the relevant instances of ($D2_{EX}$) and ($D3_{EX}$) are not both true. The only instances of (E_x) which stand a chance of avoiding both of these problems are instances in which 'X' is replaced by some complex term, like 'perception and uncontroversial testimony', or 'perception and testimony which can, in principle, be empirically verified'. But. it does not seem possible to generate these instances of (E_x) by replacing terms in (E_1) with terms that can reasonably be believed to be their synonyms; and, consequently, it seems likely that there is no form of evidentialism which is an instance of (E_x), and which discredits all deity beliefs.

Some forms of evidentialism which are generated by substitution of alleged synonyms are not instances of (E_x). But, it seems likely that these forms of evidentialism will nevertheless be vulnerable to the problems that we have just outlined. To see this, consider the following forms of evidentialism, which can all be generated by substitution of alleged synonyms:

(E_{11}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and the propositions that S reasonably believes do not confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

(E_{12}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and the propositions that S reasonably believes on the basis of perception do not confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

(E_{13}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and the propositions that S reasonably believes on the basis of perception or testimony do not confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

Each of these forms of evidentialism seems either to lead to scepticism, or to be such that, for some subject, S, the relevant instances of (D2) and (D3) do not hold.

Let us say that a form of evidentialism, E, is *too strong* to discredit all deity beliefs iff E leads to scepticism. And, let us say that E is *too weak* to discredit all deity beliefs iff E is such that, for some subject, S, the relevant instances of (D2) and (D3) don't hold. The arguments of this section strongly suggest that every form of evidentialism which is generated by substitution of alleged synonyms is either too strong or too weak to discredit all deity beliefs. In section 5.5, I will argue that there is reason to endorse a similar conclusion with respect to forms of evidentialism that are generated by complex amendments to (E_1).

5.5 Making complex amendments

An amendment to (E_1) is *complex* iff it can be generated by making an admissible amendment to a doctrine that can itself be generated by making an admissible amendment to (E_1). Some amendments of this kind seem to generate forms of evidentialism that are invulnerable to the argument of section 5.1. In what follows, I will consider a representative sample of the forms of evidentialism that can be generated by complex amendments to (E_1). By focusing on the members of this sample, I will argue that no form of evidentialism that is generated by a complex amendment is capable of discrediting all deity beliefs.

One form of evidentialism that can be generated by making a complex amendment to (E_1) is the following:

(E_{14}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not based on epistemically rational beliefs that confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

There is good reason to believe that this form of evidentialism leads to scepticism. To see this, consider the following claim (where a belief is *basic* iff it is not based on any other beliefs, and non-basic otherwise):

(UB) Every rational non-basic belief is ultimately based on basic beliefs.

When (E_{14}) is conjoined with this plausible claim, it commits us to endorsing a very strong form of foundationalism. Because of this, there is good reason to believe that (E_{14}) leads to scepticism.

To see that (E_{14}) and (UB) jointly commit us to endorsing a strong form of foundationalism, note, first of all, that (E_{14}) entails the following claim:

(F_1) With few exceptions: every rational belief, B, is based on other rational beliefs that confirm the content of B to a sufficient degree.[13]

Next, note that, when this claim is conjoined with (UB), it entails the following claim (where a belief is *properly basic* iff it is both rational and basic):

(F_2) With few exceptions: every rational belief, B, is ultimately based on properly basic beliefs that confirm the content of B to a sufficient degree.

Finally, note that the exceptional propositions to which the qualification at the start of (E_{14}) refers seem all to be self-evident propositions, like the proposition that *everything is self-identical*. From this point, we can infer that (F_2) implies the following claim:

> (F_3) With the exception of self-evident beliefs: every rational belief, B, is ultimately based on self-evident beliefs that confirm the content of B to a sufficient degree.

(F_3) states a very strong form of foundationalism. It seems clear that this form of foundationalism will lead to scepticism. For, it seems clear that hardly any of our beliefs are confirmed to a sufficient degree by beliefs that are self-evidently true. Because of this, (F_3) commits us to saying that hardly any of our beliefs are rational.

Some advocates of (E_{14}) may try to argue that (E_{14}) does not commit us to endorsing (F_3). One way in which they may try to do this is by arguing that the exceptional propositions to which the qualification at the start of (E_{14}) refers include propositions that can be known by *perception*, as well as propositions that are self-evidently true. If perceptual propositions are exceptions to the generalisation in (E_{14}), then (E_{14}) will not commit us to endorsing anything stronger than the following form of foundationalism:

> (F_4) With the exception of self-evident beliefs and perceptual beliefs: every rational belief, B, is ultimately based on self-evident beliefs and/or perceptual beliefs that confirm the content of B to a sufficient degree.

But, there is good reason to think that this form of foundationalism also leads to scepticism. To see this, recall the following sceptical hypothesis that we introduced in section 5.1:

> (TEST) Everything that I believe on the basis of testimony is a lie, which has been told to me as part of an elaborate conspiracy that is designed to prevent me from discovering the true nature of the world.

---

[13] 'Rational belief' should here be read as 'belief that is epistemically rational'.

The negation of this sceptical hypothesis is obviously entailed by each of the propositions that I believe on the basis of testimony. So, if my belief in the negation of (TEST) is epistemically irrational, then the same goes for all of the beliefs that I hold on the basis of testimony. The propositions that I believe on the basis of perception seem to be propositions that would have been true even if (TEST) had been true; and, this strongly suggests that these perceptual propositions do not confirm the negation of (TEST) to a sufficient degree. Consequently, there is reason to think that, if (F_4) is true, then my belief in the negation of (TEST) is epistemically irrational; and, this strongly suggests that (F_4) leads to scepticism.

The above comments strongly suggest that, if we want to turn (F_4) into a form of foundationalism that does not lead to scepticism, then we will have to weaken it in something like the following way:

(F_5) With the exception of self-evident beliefs, perceptual beliefs, and testimony beliefs: every rational belief, B, is ultimately based on self-evident beliefs, perceptual beliefs and/or testimony beliefs that confirm the content of B to a sufficient degree.

But, if (E_{14}) commits us to endorsing any form of foundationalism, then it commits us to endorsing a form of foundationalism that is stronger than (F_5). For, (E_{14}) tells us that there are just a few beliefs that can be both rational and basic; and, (F_5) entails that there are many beliefs of this kind. Consequently, it seems clear that, if (E_{14}) commits us to endorsing any form of foundationalism, then (E_{14}) will lead to scepticism.

It may turn out that (E_{14}) does *not* commit us to endorsing any form of foundationalism. For, it may turn out that, contrary to appearances, the following claim is untrue:

(UB) Every rational non-basic belief is ultimately based on basic beliefs.

If (UB) is not true, then it is not likely that (E_{14}) leads to scepticism. But, the falsity of (UB) also seems to commit us to saying that (E_{14}) does not discredit all deity beliefs. To see this, recall that (E_{14}) discredits all deity beliefs only if it is the case that, for every subject, S, who has deity beliefs, the following pair of claims are both true:

($D2_{E14}$) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that (E_{14}) imposes, then we can make S aware of this fact.

($D3_{E14}$) ($D2_{E14}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way.

If (UB) is not true, then it is not likely that these claims both hold with respect to every subject, S, who has deity beliefs. For, if (UB) is not true, then it will be very difficult to make such subjects aware of the fact that their deity beliefs do not satisfy the evidential requirement that (E_{14}) imposes. If (UB) is not true, then the most promising way of arguing that these beliefs do not satisfy this requirement is to argue that none of these beliefs are rational. But, if we can make a subject, S, aware of the fact that her deity beliefs are not rational before we have made her aware of the fact that her deity beliefs don't satisfy the evidential requirement that (E_{14}) imposes, then we can rebut some of S's epistemic deity beliefs in a non-evidential way. Because of this, it seems fair to conclude that, if (UB) is not true, then there are at least some subject for whom ($D2_{E14}$) and ($D3_{E14}$) are not both true. And, from this, it follows that, if (UB) is not true, then (E_{14}) does not discredit all deity beliefs.

If the above arguments succeed, then they show that (E_{14}) is either too strong, or too weak to discredit all deity beliefs. By making small alterations to these arguments, we can show that the same is true with regard to each of the following forms of evidentialism:

(E_{15}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not based on epistemically rational beliefs that confirm P to a sufficient degree, then S's belief in P does not constitute knowledge.

(E_{16}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not based on knowledge-constituting beliefs that confirm P to a sufficient degree, then S's belief in P is epistemically irrational.

(E_{17}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not based on knowledge-constituting beliefs that confirm P to a sufficient degree, then S's belief in P does not constitute knowledge.

There are other forms of evidentialism that can be generated by complex amendments to (E_1). But, none of these forms of evidentialism seem interestingly different from the forms of evidentialism that we have already discussed. All of them seem to be either too strong, or too weak to discredit all deity beliefs. Because of this, it seems likely that *no* form of evidentialism that is generated by a complex amendment is capable of discrediting all deity beliefs.

5.6 Summary

If the arguments of section 5.1 succeed, then they show that (E_1) does not discredit all deity beliefs. And, if the arguments of sections 5.2 - 5.5 succeed, then they show that the same holds with regard to all doctrines that can be generated by making an admissible amendment to (E_1). A doctrine, D, is a form of evidentialism iff D is either identical with (E_1), or can be generated by making an admissible amendment to (E_1). So, if the arguments of sections 5.1 – 5.5. succeed, then they show that there is no form of evidentialism that discredits all deity beliefs.

6. Summary and preview

In section 1, we introduced and clarified the following claim, which many people believe:

> (EBG) It is possible to undermine belief in God by arguing that the evidence for this belief is insufficient.

Sections 2–5 then outlined an objection to this claim, which we called *the anti-sceptical objection*. In the next four chapters, I will outline and solve two serious problems which are faced by advocates of the anti-sceptical objection. By solving these problems, I hope to show that the objection is worthy of further attention.

Chapter Two

A problem for the anti-sceptical objection

In the last chapter, we outlined an objection to the following claim:

> (EBG) It is possible to undermine belief in God by arguing that the evidence for this belief is insufficient.

We called this objection *the anti-sceptical objection*, or *the ASO*. The aim of this chapter is to show that there is a serious problem which is faced by advocates of the ASO. In chapter three, we will show that the ASO can be made invulnerable to this problem.

This chapter is divided into six sections. Section 1 describes a situation in which a woman, Mary, gives her husband, John, a conclusive epistemic reason to abandon a certain belief. It then outlines, and briefly defends a knowledge-based explanation of the way in which Mary undermines John's belief. Sections 2–5 improve section 1's defence of this knowledge-based explanation by arguing for two principles about knowledge which jointly commit us to saying that this explanation is correct. And section 6 uses the explanation to generate a problem for the anti-sceptical objection by arguing (a) that, if the explanation is correct, then there is a form of evidentialism that discredits many defective beliefs, and (b) that, by appealing to this form of evidentialism, we can show that the second stage of the anti-sceptical objection fails.

1. The job interview

It is Monday afternoon, and John has just come out of a job interview. The interview has gone well, so John's spirits are high. At 4pm, he meets his wife Mary for a cup of coffee, and she asks him how the interview went. He responds that it went very well, and tells Mary that he is going to get the job. Mary knows from experience that John is prone to wishful thinking on this topic, so she asks him why he thinks that he is going to get the job. The following exchange ensues:

> *John:* Why do I think that I'm going to get the job? Because the interview went really well. I got on well with both of the interviewers, and they both seemed really pleased with my answers to their questions.
>
> *Mary:* Did they actually tell you that you were going to get the job?
>
> *John:* Well, no. But, like I said, I got on really well with them. And, I could tell that they were impressed with my performance.
>
> *Mary:* It's good that you got on well with them. And, it's good that they were impressed. But, none of that establishes that you're going to get the job. There might be several other candidates who got on just as well with them, and who impressed them just as much. And, some of those candidates might be better qualified than you are...

If the situation here is normal, then Mary's questions and comments will give John a conclusive epistemic reason to abandon his belief about the job (henceforth: his *job belief*). The question that I want to address is: *why* do these questions and comments give John a conclusive epistemic reason to abandon his job belief? In section 1.1, I will sketch an answer to this question. Sections 1.2–1.6 will refine this answer by appealing to other situations like the one just described.

1.1 The irrationality explanation

One way of giving a subject, S, a conclusive epistemic reason to abandon a belief, B, is to make S aware of the fact that B is epistemically irrational. It is natural to suppose that Mary gives John reason to abandon his job belief in this way. For, it seems clear that Mary's questions and comments *do* make John aware of the fact that his job belief is epistemically irrational. And, it is hard to see how we can explain the undermining effect of Mary's questions and comments without appealing to something like this fact.

How does Mary make John aware of the fact that his job belief is epistemically irrational? Roughly: by asking him *why* he has this belief. More exactly: by asking him for an explanation of this belief, and then arguing that this explanation does not establish that his job belief is true. Mary's questions and comments seem to make it clear to John that the following claims are both true with respect to him:

(B_j) I believe that *I will get the job* because my interview went well.

($\neg E_j$) The fact that my interview went well does not establish that I will get the job.

Consequently, it is natural to give something like the following explanation of the effect of Mary's questions and comments:

Mary's questions and comments give John a conclusive epistemic reason to abandon his job belief because they make him aware of the fact that (B_j) and ($\neg E_j$) are both true with respect to him. John is already aware of the fact that, if (B_j) and ($\neg E_j$) are both true with respect to him, then the following claim is also true with respect to him:

(I_j) My belief that *I will get the job* is epistemically irrational.

So, by making John aware of the fact that (B_j) and ($\neg E_j$) are true with respect to him, Mary makes him aware of the fact that his job belief is epistemically irrational. And, by making him aware of this fact, she gives him a conclusive epistemic reason to abandon his job belief.

Let us refer to this explanation as *the irrationality explanation.* In the next two subsections, we will argue for a modification to the irrationality explanation. This modification will be motivated by reflection on another case in which Mary undermines one of John's beliefs. Section 1.2 will describe this case, and section 1.3 will outline the modification that it prompts.

1.2 The broken thermometer[1]

It is Tuesday morning, and John and Mary have woken up to discover that their son, Joe, is ill. John is worried that Joe might have the 'flu, so he uses Mary's thermometer to take Joe's temperature. After leaving the thermometer in Joe's mouth for several minutes, John examines it, and is relieved to see that the reading on the thermometer is normal – i.e. that it reads 98.6 degrees F. He goes downstairs to tell Mary this news, and the following exchange ensues:

> *John*: Well, it looks like Joe doesn't have the flu. His temperature is normal.
>
> *Mary*: Why do you think that his temperature is normal?
>
> *John*: Because your thermometer says that it's only 98.6.
>
> *Mary:* That doesn't establish that his temperature is normal. My thermometer isn't working properly. It's working okay for temperatures below 98 degrees, but it isn't registering any temperature that is higher than that...

If the situation here is normal, then Mary's questions and comments will give John a conclusive epistemic reason to abandon his belief about Joe's temperature (henceforth: his *temperature belief*). Once again, I am interested in *why* these questions and comments have this effect on John. It seems fairly clear that they undermine his temperature belief in the same way as Mary's earlier questions and comments undermined his job belief. So, if the irrationality explanation is correct, then the following, parallel explanation should also be correct:

Mary's questions and comments give John a conclusive epistemic reason to abandon his temperature belief because they make him aware of the fact that the following claims are true with respect to him:

[1] The example given in this section is derived from an example given by Dretske. See Dretske 1971:2.

(B_1) I believe that *Joe's temperature is normal* because Mary's thermometer says that it is normal.

($\neg E_1$) The fact that Mary's thermometer says that Joe's temperature is normal does not establish that Joe's temperature is normal.

John is already aware of the fact that, if (B_1) and ($\neg E_1$) are true with respect to him, then the following claim is also true with respect to him:

(I_1) My belief that *Joe's temperature is normal* is epistemically irrational.

So, by making John aware of the fact that (B_1) and ($\neg E_1$) are true with respect to him, Mary makes him aware of the fact that his temperature belief is epistemically irrational. And, by making him aware of this fact, she gives him a conclusive epistemic reason to abandon his temperature belief.

In what follows, we will refer to this explanation as *the second irrationality explanation*. There is good reason to think that the second irrationality explanation is false. For, this explanation entails that, in the situation just described, John's temperature belief is epistemically irrational. And, it seems clear that, in this situation, John's temperature belief is *not* epistemically irrational. If John had been aware of the problems with the thermometer, or if he had had good reason to suspect that there were such problems, then it might have been irrational for him to believe, on the basis of the thermometer reading, that Joe's temperature was normal. But, there is nothing in the above description to indicate that John was, or should have been aware of these problems. Consequently, we can conclude that the second irrationality explanation is false. And, this gives us good reason to believe that the original irrationality explanation is also false.

1.3 The lack-of-knowledge explanation

The arguments of the last subsection show that we need to modify the irrationality explanation. One natural way of modifying it is to rewrite it in the following way:

Mary's questions and comments give John a conclusive epistemic reason to abandon his job belief because they make him aware of the fact that the following claims are both true with respect to him:

(B_j) I believe that *I will get the job* because my interview went well.

($\neg E_j$) The fact that my interview went well does not establish that I will get the job.

John is already aware of the fact that, if (B_j) and ($\neg E_j$) are true with respect to him, then the following claim is also true with respect to him:

($\neg K_j$) I do not know that I will get the job.

So, by making John aware of the fact that (B_j) and ($\neg E_j$) are true with respect to him, Mary makes him aware of the fact that his job belief does not constitute knowledge. And, by making him aware of this fact, she gives him a conclusive epistemic reason to abandon his job belief.

Let us refer to this explanation as the *lack of knowledge*, or *LOK* explanation. It seems clear that, in the case of the broken thermometer, Mary *does* make John aware of the fact that his temperature belief does not constitute knowledge; so, the LOK explanation is not vulnerable to the objection that we outlined in section 1.2. Some may think that the explanation is nevertheless vulnerable to other objections. In sections 1.4–1.6 , I will try to undermine this thought by rebutting two natural objections to the LOK explanation.

1.4 Irrationality and the LOK explanation

One immediate problem with the LOK explanation can be generated by focusing on one of the most attractive features of the irrationality explanation. It is very natural to suppose that, in the case of the job interview, Mary makes John aware of the fact that his job belief is epistemically irrational. And, it is equally natural to say that Mary makes John aware of this fact by making him aware of the fact that the following claims are both true with respect to him:

(B_j) I believe that *I will get the job* because my interview went well.

($\neg E_j$) The fact that my interview went well does not establish that I will get the job.

The irrationality explanation accommodates both of these intuitions nicely. But, the LOK explanation does not appear to accommodate either. Some may think that this gives us reason to prefer the irrationality explanation. In what follows, I will attack this thought, by showing that the LOK explanation *can* accommodate both of the intuitions just described.

When John tells Mary that he is going to get the job, he is not aware of the fact that (B_j) and ($\neg E_j$) are both true with respect to him. Nevertheless, it seems clear that he *ought* to be aware of this fact. For, he knows that his job belief is explained by the fact that his interview went well. And, he knows – or at least, ought to know – that this fact does not establish that his job belief is true.

If the LOK explanation is correct, and John ought to be aware of the fact that (B_j) and ($\neg E_j$) are true with respect to him, then John ought to be aware of the fact that his job belief does not constitute knowledge. For, the LOK explanation tells us that John is already aware of the fact that the following claim is true with respect to him:

(BEK_j) If (B_j) and ($\neg E_j$) are true with respect to me, then ($\neg K_j$) is true with respect to me.

And, if John is aware of this fact, and of the fact that the antecedent of (BEK_j) is true with respect to him, then he should also be aware of the fact that the consequent of (BEK_j) is true with respect to him.

If the LOK explanation is correct, and John ought to be aware of the fact that his job belief does not constitute knowledge, then John ought to have a conclusive epistemic reason to abandon his job belief. And, if John ought to have a conclusive epistemic reason to abandon this belief, then it is surely epistemically irrational for him to have this belief (given that the 'ought' that we are using here is the 'ought' of epistemic come up. That ticket wasn't going to win the lottery, so there was no point in keeping hold of it.

Mary: Why do you think that your numbers aren't going to come up?

John: Because the chances of them coming up are ridiculously small. There's more chance that I'll be struck by lightning tomorrow than there is of my numbers coming up tonight.

Mary: That may be true; but, even if it is, it doesn't *establish* that your numbers aren't going to come up tonight. If it did, then similar considerations would establish the same thing with respect to *every* set of numbers that could come up. One of those sets of numbers *will* come up tonight, even though the chance of it doing so is extremely small. So, you can't conclude that your numbers won't come up just because the chance of them doing so is extremely small.

If the situation here is normal, then Mary's questions and comments will make John aware of the fact that the following claims are true with respect to him:

(B_n) I believe that *my numbers will not come up* because the chance of them coming up is extremely small.

($\neg E_n$) The fact that *the chance of my numbers coming up is extremely small* does not establish that my numbers will not come up.

But, it is not clear that they will give John a conclusive epistemic reason to stop believing that his numbers will not come up. For, it seems clear that John can reasonably remain very confident of the claim that his numbers will not come up, after being made aware of the fact that (B_n) and ($\neg E_n$) are true with respect to him. And, if he can reasonably remain very confident of this claim, then it's natural to say that he can also reasonably continue to *believe* it.

The LOK explanation commits us to saying that John *cannot* reasonably continue to believe that his numbers will not come up. So, the above comments threaten to show that the LOK explanation is false. In the next subsection, I will defend the LOK explanation against this threat. I will then mount a positive defence of the explanation, by arguing for two principles about knowledge which jointly commit us to saying that the explanation is correct.

1.6 Active versus passive belief

To defend the LOK explanation against the objection outlined in the last section, we need to return to the distinction between active and passive belief that we made in chapter one. As in chapter one, we can make this distinction by focusing on the following example, which is due to Alvin Plantinga:

> 'Consider a Christian beset by doubts. He has a hard time believing certain crucial Christian claims – perhaps the teaching that God was in Christ, reconciling the world to himself. Upon calling that belief to mind, he finds it cold, lifeless, without warmth or attractiveness. Nonetheless, he is committed to this belief; it is his position; if you ask him what he thinks about it, he will unhesitatingly endorse it. He has, so to speak, thrown in his lot with it.'
>
> (Plantinga 1983:37).

To actively believe the proposition that p is to be *committed* to this proposition, in the way that Plantinga's Christian is committed to the doctrine that Plantinga mentions. And, to passively believe it is to have the kind of confidence in it that Plantinga's Christian lacks, with respect to the proposition that Plantinga describes. By appealing to the distinction between active and passive belief, we can distinguish two readings of the LOK explanation. On the first reading, the explanation is meant to describe the way in which Mary undermines John's *active* job belief, and on the second, it is meant to describe the way in which she undermines his *passive* job belief. If the explanation is read in the second way, then it may well be vulnerable to the lottery objection that we outlined in section 1.5. But, if it is read in the first way, then it will not be vulnerable to this objection, as we will now show.

In the lottery case that we described in section 1.5, Mary makes John aware of the fact that the following claims are true with respect to him:

(B_n) I believe that *my numbers will not come up* because the chance of them coming up is extremely small.

($\neg E_n$) The fact that *the chance of my numbers coming up is extremely small* does not establish that my numbers will not come up.

The core claim of the lottery-based objection is that it is reasonable for John to believe that his numbers will not come up even after he has been made aware of the fact that these claims are true with respect to him. If this core claim is read as a claim about passive belief, then it is plausible – for, it seems reasonable for John to be very confident that his numbers will not come up, even after he has learned that (B_l) and ($\neg E_l$) are true with respect to him. But, if it is read as a claim about active belief, then it is implausible – for, when John learns that (B_l) and ($\neg E_l$) are true with respect to him, he seems to be given a conclusive epistemic reason to refrain from actively *asserting* or *endorsing* the claim that his numbers will not come up.

The above comments suggest that the LOK explanation will be vulnerable to the lottery-based objection only if it is read as a description of the way in which Mary undermines John's *passive* job belief. In what follows, we will adopt an alternative reading, on which the aim is to describe the way in which Mary undermines John's *active* job belief. By adopting this reading, we seem to make the LOK explanation invulnerable to the lottery-based objection. In sections 2–5, we will construct a positive case for the LOK explanation, by arguing for two principles which jointly commit us to saying that the explanation is correct.

2. Two principles about knowledge

In the next three sections, I will defend two principles about knowledge. The first of these principles says something about the relationship between knowledge and establishment, and the second says something about the relationship between knowledge and reasons for belief. In this section, I will state the principles, and will argue that acceptance of them commits us to accepting the LOK explanation. Later sections will argue for both of the principles, and will thereby generate an argument for the LOK explanation.

The first principle that I will defend says something about the relationship between knowledge and establishment. If we appeal to the following schematic sentences:

(B) (S believes that p) because q.
(E) The fact that q establishes that p.
(K) S knows that p.

then we can state the principle in the following way:

(BEK) We are all aware of the fact that, if an instance of (B) is true, and the corresponding instance of (E) is untrue, then the corresponding instance of (K) is also untrue.

The second principle that I will defend says something about the relationship between knowledge and active belief. If 'S' and 'P' range over all subjects and propositions, then the principle can be stated in the following way:

(KAB) If S actively believes P, and S is aware of the fact that she does not know P, then S has a conclusive epistemic reason to stop actively believing P.

It should be fairly clear that, if we accept both (BEK) and (KAB), then we will be committed to endorsing the LOK explanation. For, it is clear that, if (BEK) is true, then the following claim about the job interview example is also true:

(J_1) John is aware of the fact that, if (B_j) and ($\neg E_j$) are both true with respect to him, then ($\neg K_j$) is also true with respect to him.

And, it is equally clear that, if (KAB) is true, then this claim about the job interview example is also true:

(J_2) If John is aware of the fact that ($\neg K_j$) is true with respect to him, then John has a conclusive epistemic reason to stop actively believing that he will get the job.

Claims (J_1) and (J_2) are the only controversial claims that advocates of the LOK explanation are committed to making. All of the other claims that they make should be accepted without question. So, if (BEK) and (KAB) are true, then we ought to endorse the LOK explanation. In the next three sections, I will argue that (BEK) and (KAB) *are* both true.

3. Arguing for (BEK)

In this section, I will give an argument for (BEK). This argument will be given in three stages. The first stage (sections 3.1–3.3) will outline three puzzles about establishment, and will argue that the best way of solving these puzzles is to appeal to the following psychological generalisation:[2]

(CE_j) When we judge that an instance of the following claim is untrue:

(C) If it was not the case that p, then it wouldn't be the case that q

we tend also to judge that the corresponding instance of (E) is untrue.

The second stage (section 3.4) will use the results of the first stage to argue that, if (BEK) is true, then the following psychological generalisation should also be true:

(BCK_j) When we judge that an instance of (B) is true, and that the corresponding instance of (C) is untrue, we tend also to judge that the corresponding instance of (K) is untrue.

And, the third stage (sections 3.5–3.7) will argue that this psychological generalisation *is* true. If all three stages of this argument succeed, then they will provide us with a strong reason for believing (BEK). In section 4, we will give another reason for believing (BEK), and in section 4, we will do likewise with respect to (KAB).

3.1 A puzzle about lotteries[3]

In my daily paper, there is a list of football scores which says that, in yesterday's game, Arsenal beat Manchester United. This fact about my paper – which we will call *the newspaper fact* – provides very strong evidence for the claim that Arsenal *did* beat Manchester United. There may be other facts that provide stronger evidence for this claim. But, in spite of this, it seems clear that the following instance of (E) is true:

(E_n) The newspaper fact *establishes* that Arsenal beat United.

In last Saturday's National Lottery, the chance of my ticket winning the jackpot was lower than one in ten million. This fact about the lottery – which we will call *the chance fact* – provides very strong evidence for the claim that my ticket did not win the jackpot. The evidence that it provides for this claim seems at least as strong as the evidence that the newspaper fact provides for the claim that Arsenal beat United. But, in spite of this, it seems clear that the following instance of (E) is *not* true:

(E_l) The chance fact establishes that my ticket did not win the jackpot.

Why does (E_l) seem untrue, when the chance fact provides at least as much evidence for the claim that *my ticket did not win* as the newspaper fact provides for the claim that *Arsenal beat United*? One natural way of explaining this is to appeal to (CE_j). Corresponding to (E_n) and (E_l) are the following instances of (C):

---

[2] The argument of these sections is inspired by a similar argument in DeRose 1995.
[3] This puzzle is based on a similar puzzle described in Harman 1968. The solution is similar to a

(C_n) If Arsenal had not beaten United, then my newspaper would not have said that Arsenal beat United.

(C_l) If my ticket had won the jackpot, then the chance of it winning the jackpot would not have been lower than one in ten million.

Although (C_n) seems true, (C_l) seems untrue. (CE_j) tells us that, if we judge that an instance of (C) is untrue, then we tend also to judge that the corresponding instance of (E) is untrue - so, (CE_j) can explain why (E_l) seems untrue even though (E_n) does not. The fact that (CE_j) can explain this is a reason to believe (CE_j). In the next two subsections, we will outline two similar reasons for believing (CE_j).

3.2 A puzzle about zebras[4]

During a visit to my local zoo, I stop outside the zebra enclosure. Inside the enclosure, there are striped animals which look to me like zebras, and a sign which says that these animals *are* zebras. This fact about the contents of the enclosure - which we will call *the contents fact* - provides very strong evidence for the claim that the animals in the enclosure are zebras. There may be other facts that provide stronger evidence for this claim; but, it nevertheless seems clear that the following claim is true:

(E_z) The contents fact establishes that the animals in the enclosure are zebras.

Although (E_z) seems true, the following claim seems untrue:

(E_m) The contents fact establishes that the animals in the enclosure are not mules that have been cleverly disguised as zebras by the people who own the zoo (henceforth: that the animals in the enclosure are not *cleverly disguised mules*).

---

solution given in DeRose 1996.

[4] This puzzle is based on a similar puzzle described in Dretske 1970. The solution is similar to the

Why does this claim seem untrue, when the contents fact seems to establish that the animals in the enclosure are *zebras*, and hence not cleverly disguised mules? Again, we can explain this by appealing to (CE_j). Corresponding to (E_z) and (E_m) are the following instances of (C):

(C_z) If the animals in the enclosure were not zebras, then it would not be the case that they look to me like zebras, and are referred to as zebras by the sign in their enclosure.

(C_m) If the animals in the enclosure were cleverly disguised mules, then it would not be the case that they look to me like zebras, and are referred to as zebras by the sign in their enclosure.

Although (C_z) seems true, (C_m) seems untrue.[5] So, (CE_j) can explain why (E_m) seems untrue even though (E_z) seems true. The fact that it can explain this is another reason for believing (CE_j). In the next subsection, we will outline a third reason of this kind; and, in later sections, we will argue that, if (CE_j) is true, then there is good reason to believe (BEK).

3.3 A puzzle about contrasts[6]

My friend Jim has just bought a new suit, and is wearing this suit as he walks down the High Street of our local town. As he walks past the window of the coffee shop where I am sitting, I see him, and see that he is wearing a suit that looks new. The fact that Jim is wearing a suit that looks new – which we will call *the suit fact* - strongly supports the claim that Jim has just bought a new suit. There may be other facts that support this claim to a higher degree; but, even if there are, it seems right to say that the following claim is true:

(E_b) The suit fact establishes that Jim has just bought a new suit.

solution given in DeRose 1995.

[5] After reflection on (C_m), (C_z) can sometimes seem untrue. But, when it does, (E_z) also seems untrue: so this fact supports (CE_j), rather than threatening it.

[6] This puzzle is inspired by the discussion of knowledge and contrasts on pp.1021ff. of Dretske 1970.

Although (E_b) seems true, the following claim seems untrue:

(E_s) The suit fact establishes that Jim has just *bought*, rather than stolen, a new suit.

Why does (E_s) seem untrue, when the suit fact seems to establish that Jim has just bought a new suit? Again, we can explain this by appealing to (CE_j). Corresponding to (E_b) and (E_s) are the following instances of (C):

(C_b) If Jim had not just bought a new suit, then Jim would not be wearing a suit that looks new.

(C_s) If Jim had just stolen, rather than bought some a new suit, then Jim would not be wearing a suit that looks new.

Although (C_b) seems true, (C_s) seems untrue.[7] So, (CE_j) can explain why (E_s) seems untrue even though (E_b) seems true.

3.4 (CE_j) and (BEK)

The last three subsections have provided us with three good reasons for believing (CE_j). The next four subsections will argue that, if (CE_j) is true, then we have good reason to believe (BEK).

According to (BEK), we are all aware of the fact that, if an instance of (B) is true, and the corresponding instance of (E) is untrue, then the corresponding instance of (K) is also untrue. It seems clear that, if we are all aware of this fact, then the following psychological generalisation will be true:

(BEK_j) When we judge that an instance of (B) is true, and that the corresponding instance of (E) is untrue, we tend to judge that the corresponding instance of (K) is also untrue.

---

[7] After reflection on (C_s), (C_b) can sometimes seem untrue. But, when it does, (E_b) also seems untrue: so, this fact supports (CE_j), rather than threatening it.

When this generalisation is conjoined with (CE_j), it commits us to endorsing another generalisation, which can be stated in the following way:

> (BCK_j) When we judge that an instance of (B) is true, and that the corresponding instance of (C) is untrue, we tend to judge that the corresponding instance of (K) is also untrue.

So, if (BEK) is true, then (BCK_j) should also be true. In the next three sections, I will argue that (BCK_j) *is* true. I will argue for this by arguing that, by appealing to (BCK_j), we can solve three puzzles about knowledge that are exactly parallel to the puzzles that we outlined in sections 3.1–3.3. If my arguments succeed, then they will give us reason to believe (BEK). Sections 4 will describe another reason for believing (BEK).

3.5 *Another puzzle about lotteries*

In my daily paper, there is a list of football scores which says that, in yesterday's game, Arsenal beat Manchester United. This fact about my paper – which we are calling *the newspaper fact* – provides very strong evidence for the claim that Arsenal *did* beat Manchester United. There may be other facts that provide stronger evidence for this claim. But, in spite of this, it seems clear that, if I believe, on the basis of the newspaper fact, that Arsenal beat United, then the following claim is true:

> (K_n) I know that Arsenal beat United.

In last Saturday's National Lottery, the chance of my ticket winning the jackpot was lower than one in ten million. This fact about the lottery – which we are calling *the chance fact* – provides very strong evidence for the claim that my ticket did not win the jackpot. The evidence that it provides for this claim seems at least as strong as the evidence that the newspaper fact provides for the claim that Arsenal beat United. But, in spite of this, it seems clear that, if I believe, on the basis of the chance fact, that my ticket did not win the jackpot, then the following claim is *not* true:

(K_l) I know that my ticket did not win the jackpot.

Why does this seem clear, when the chance fact provides as much evidence for the claim that *my ticket did not win* as the newspaper fact provides for the claim that *Arsenal beat United*? One natural way of explaining this is to appeal to (BCK_j). If I believe, on the basis of the newspaper fact, that Arsenal beat United, then the following claim will be true:

(B_n) I believe that *Arsenal beat United* because my newspaper says that Arsenal beat United.

And, if I believe, on the basis of the chance fact, that I did not win the jackpot, then this claim will be true:

(B_l) I believe that *my ticket did not win the jackpot* because the chance of my ticket winning the jackpot was lower than one in ten million.

Corresponding to these two instances of (B) are the following instances of (C):

(C_n) If Arsenal had not beaten United, then my newspaper would not have said that Arsenal beat United.

(C_l) If my ticket had won the jackpot, then the chance of it winning the jackpot would not have been lower than one in ten million.

Although (C_n) seems true, (C_l) seems untrue. (BCK_j) tells us that, if we judge that an instance of (B) is true, and we judge that the corresponding instance of (C) is untrue, then we tend also to judge that the corresponding instance of (K) is untrue. So, (BCK_j) can explain why it is that, in the above-described circumstances, (K_l) seems untrue even though (K_n) seems true. The fact that (BCK_j) can explain this is a reason to believe (BCK_j).

3.6 *Another puzzle about zebras*

During a visit to my local zoo, I stop outside the zebra enclosure. Inside the enclosure, there are striped animals which look to me like zebras, and a sign which says that these animals *are* zebras. This fact about the contents of the enclosure - which we are calling *the contents fact* - provides very strong evidence for the claim that the animals in the enclosure are zebras. There may be other facts that provide stronger evidence for this claim; but, it nevertheless seems clear that, if I believe, on the basis of the contents fact, that the animals in the enclosure are zebras, then the following claim is true:

(K_z) I know that the animals in the enclosure are zebras.

Although this seems clear, it also seems clear that, if I believe, on the basis of the contents fact, that the animals in the cage are not *cleverly disguised mules*, then the following claim is *not* true:

(K_m) I know that the animals in the enclosure are not cleverly disguised mules

Why does this seem clear, when I seem able to know, on the basis of the contents fact, that the animals in the enclosure are *zebras* - and, hence, not cleverly disguised mules? Again, we can explain this by appealing to (BCK_j). If I believe, on the basis of the contents fact, that the animals in the enclosure are zebras, then the following claim will be true:

(B_z) I believe that *the animals in the enclosure are zebras* because the animals in the enclosure look to me like zebras, and are referred to as zebras by the sign in the enclosure.

And, if I believe, on the basis of the contents fact, that the animals in the enclosure are not cleverly painted mules, then this claim will be true:

(B_m) I believe that *the animals in the enclosure are not cleverly disguised mules* because the animals in the enclosure look to me like zebras, and are referred to as zebras by the sign in the enclosure.

Corresponding to these two instances of (B) are the following instances of (C):

(C_z) If the animals in the enclosure were not zebras, then it would not be the case that they look to me like zebras, and are referred to as zebras by the sign in the enclosure.

(C_m) If the animals in the enclosure were cleverly disguised mules, then it would not be the case that they look to me like zebras, and are referred to as zebras by the sign in the enclosure.

Although (C_z) seems true, (C_m) seems untrue.[8] So, (BCK_j) can explain why it is that, in the above-described circumstances, (K_m) seems untrue even though (K_z) seems true. The fact that (BCK_j) can explain this is another reason for believing (BCK_j). In the next subsection, I will outline a third reason of this kind.

3.7 Another puzzle about contrasts

My friend Jim has just bought a new suit, and is wearing this suit as he walks down the High Street of our local town. As he walks past the window of the coffee shop where I am sitting, I see him, and see that he is wearing a suit that looks new. The fact that Jim is wearing a suit that looks new – which we are calling *the suit fact* - strongly supports the claim that Jim has just bought a new suit. There may be other facts that support this claim to a higher degree; but, even if there are, it seems right to say that, if I believe, on the basis of the suit fact, that Jim has just bought a new suit, then the following claim is true:

(K_b) I know that Jim has just bought a new suit.

[8] After reflection on (C_m), (C_z) can sometimes seem untrue. But, when it does, (K_z) also seems untrue: so, this fact supports (CE_j), rather than threatening it.

Although this seems clear, it also seems clear that, if I believe, on the basis of the suit fact, that Jim has *bought*, rather than stolen a new suit, then the following claim is *not* true:

(K_s) I know that Jim has just *bought*, rather than stolen, a new suit.

Why does this seem clear, when I seem able to know, on the basis of the suit fact, that Jim has just bought a new suit? Again, we can explain this by appealing to (BCK_j). If I believe, on the basis of the suit fact, that Jim has just bought a new suit, then the following claim is true:

(B_b) I believe that *Jim has just bought a new suit* because Jim is wearing a suit that looks new.

And, if I believe, on the basis of the suit fact, that Jim has just *bought*, rather than stolen a new suit, then this claim is true:

(B_s) I believe that *Jim has just bought, rather than stolen a new suit* because Jim is wearing a suit that looks new.

Corresponding to (B_b) and (B_s) are the following instances of (C):

(C_b) If Jim had not just bought a new suit, then Jim would not be wearing a suit that looks new.

(C_s) If Jim had just stolen, rather than bought a new suit, then Jim would not be wearing a suit that looks new.

Although (C_b) seems true, (C_s) seems untrue.[9] So, (BCK_j) can explain why it is that, in the above-described circumstances, (K_b) seems untrue even though (K_s) seems true. The fact that (BCK_j) can explain this provides us with another reason for

[9] After reflection on (C_s), (C_b) can sometimes seem untrue. But, when it does, (K_b) also seems untrue: so, this fact supports (CE_j), rather than threatening it.

endorsing (BCK_j). Consequently, we now have good reason to believe (BCK_j).

4. Another argument for (BEK)

The arguments of the last section have provided us with good reasons for believing both (CE_j) and (BCK_j). By doing so, they have also given us reason to believe (BEK). In this section, I will use the arguments of the last section to generate another argument for (BEK). Section 4.1 will argue that the best way of explaining the truth of (CE_j) is to endorse the following contextualist claim about establishment[10]:

> (CCE) When it is asserted (or implied) that an instance of (E) is true (or untrue), the standards for establishment tend to be raised in such a way as to ensure that, if that instance of (E) is true, then the corresponding instance of (C) is also true.

Section 4.2 will use the results of section 4.1 to argue that, if (BEK) is true, then the following contextualist claim about knowledge should also be true:

> (CCK) When it is asserted (or implied) that an instance of (B) is true and that the corresponding instance of (K) is true, the standards for knowledge tend to be raised in such a way as to ensure that, if those instances of (B) and (K) are true then the corresponding instance of (C) is also true.

And, section 4.3 will argue that this contextualist claim about knowledge *is* true, and that there is consequently another reason for endorsing (BEK).[11]

---

[10] I call this a contextualist claim because it entails that the truth-conditions of claims about establishment vary with the context in which those claims are made. Both this claim and (CCK) resemble contextualist claims defended in DeRose 1995.

[11] The arguments of sections 4.1 and 4.3 are inspired by similar arguments in sections 9-13 of DeRose 1995.

4.1 *Explaining (CE_j)*

If (CE_j) is true, then it is natural to ask *why* it is true. Why do we tend to judge that instances of (E) are untrue, when we judge that the corresponding instances of (C) are untrue? The most straightforward way of explaining this is to endorse the following claim, which closely resembles (CE_j):

> (EC) We are all aware of the fact that, if an instance of (E) is true, then the corresponding instance of (C) is true.

If (EC) is true, then we are presumably also aware of the fact that, if an instance of (C) is untrue, then the corresponding instance of (E) is untrue. And, if we are aware of this fact, then it is not surprising that we tend to judge that instances of (E) are untrue, when we judge that the corresponding instances of (C) are untrue.

The problem with this explanation is that there are good reasons for thinking that (EC) is false. One good way of seeing this is to consider the relationship between (EC) and the following plausible principle about establishment:

> (PE) If the fact that p establishes that q, and the claim that q entails that r, then the fact that p establishes that r.

When (EC) is conjoined with this principle, it entails that many of our everyday claims about establishment are false. One good way of seeing this is to focus on some of the claims about establishment that were made in sections 3.1–3.3. In the situation described in section 3.2, it seemed clear that the following claim was true:

> (E_z) The contents fact establishes that the animals in the enclosure are zebras.

But, if (EC) and (PE) are both true, then (E_z) is false. To see this, note first of all that, if (PE) is true, then (E_z) implies the following claim:

(E_m) The contents fact establishes that the animals in the enclosure are not cleverly disguised mules.

Next, note that, if (EC) is true, then (E_m) implies the following claim:

(C_m) If the animals in the enclosure were cleverly disguised mules, then it would not be the case that they look to me like zebras, and are referred to as zebras by the sign in the enclosure.

Finally: note that, in the situation described in section 3.2, (C_m) is clearly false. From this, it follows that, if (EC) and (PE) are both true, then, in the situation described in section 3.2, (E_z) is false. But, it seems clear that, in the situation described in section 3.2, (E_z) is *not* false. And, for this reason, it seems clear that, if we endorse (PE), then we should reject (EC).

The most natural way of responding to this attack on (EC) is to try to show that (PE) is false. But, it seems unlikely that a response of this kind will succeed. For, by responding to the attack in this way, we commit ourselves to endorsing 'abominable conjunctions'[12] like the following:

(AC_E) The contents fact establishes that the animals in the enclosure are zebras, but it does not establish that they are not cleverly painted mules.

And, most of us are very strongly inclined to refrain from endorsing conjunctions of this kind.

Consequently, it is natural to say that (EC) is false. But, if we say this, then we need another explanation for the truth of (CE_j). One natural way of supplying this explanation is to appeal to the following contextualist claim about establishment:

[12] The phrase 'abominable conjunction' is taken from DeRose 1995, where it is used to refer to certain conjunctions that are very similar to (AC_E).

(CCE) When it is asserted (or implied) that an instance of (E) is true (or untrue), the standards for establishment tend to be raised in such a way as to ensure that, if that instance of (E) is true, then the corresponding instance of (C) is also true.

This claim entails that, if we assert that an instance of (E) is true when the corresponding instance of (C) is untrue, then the standards for establishment will tend to be raised in such a way as to make our assertion false. It also entails that, if we assert that an instance of (E) is untrue when the corresponding instance of (C) is untrue, then the standards for establishment will tend to be raised in such a way as to make our assertion *true*. If the standards for establishment vary in this way, then it is not surprising that we tend to judge that instances of (E) are untrue, when we make the same judgement about corresponding instances of (C). So, if (CCE) is true, then it provides us with an explanation for the truth of (CE_j).[13]

Does this contextualist explanation suffer from the same problems as the explanation that appealed to (EC)? On reflection, it is clear that it does not. For, unlike the (EC) explanation, the contextualist explanation does not commit us to saying that (E_z) is false. All that it commits us to saying is that, when it is asserted (or implied) that (E_m) is true (or not true), the standards for establishment tend to be raised in such a way as to make (E_z) false. And, this claim is fairly plausible. In the situation described in section 3.2, it seems right to say that, once (E_m) (or its negation) has been asserted (or implied) we can no longer truly assert (E_z). What seems wrong is the claim that we can *never* truly assert (E_z). The (EC)-explanation commits us to making this claim, but the contextualist explanation does not. Consequently, we should opt for the contextualist explanation, by endorsing (CCE) rather than (EC).

4.2 (CCE) and (BEK)

According to (BEK), we are all aware of the fact that, if an instance of (B) is true, and the corresponding instance of (E) is untrue, then the corresponding instance of

(K) is also untrue. If we are all aware of this fact, then it is plausible to suppose that the following claim is true:

(ABK) When it is asserted, or implied that an instance of (B) is true, and that the corresponding instance of (K) is true, it is implied that the corresponding instance of (E) is also true.

When this claim about assertion is conjoined with the contextualist claim, (CCE), that we defended in the last section, it entails another contextualist claim about establishment:

(CCE_2)When it is asserted, or implied, that an instance of (B) is true, and that the corresponding instance of (K) is true, the standards for establishment tend to be raised in such a way as to ensure that, if the corresponding instance of (E) is true, then the corresponding instance of (C) is also true.

And, when this claim is conjoined with the following conditional, which is entailed by (BEK):

(BKE) If an instance of (B) is true, and the corresponding instance of (K) is true, then the corresponding instance of (E) is also true.

it entails the following contextualist claim about knowledge:

(CCK) When it is asserted (or implied) that an instance of (B) is true and that the corresponding instance of (K) is true, the standards for knowledge tend to be raised in such a way as to ensure that, if those instances of (B) and (K) are true then the corresponding instance of (C) is true.

Consequently, we can conclude that, if (BEK) is true, then (CCK) should also be true. In the next section, we will argue that (CCK) *is* true, by arguing that we need to

[13] Cf. DeRose 1995, section 13.

appeal to it in order to explain the truth of (BCK_j). If the argument of this section succeeds, then it will provide us with another reason for believing (BEK). This reason seems at least as strong as the reason for believing (BEK) that we gave in section 3.

4.3 Explaining (BCK_j)

If (BCK_j) is true, then it is natural to ask *why* it is true. Why is it that, when we judge that an instance of (B) is true, and that the corresponding instance of (C) is untrue, we tend also to judge that the corresponding instance of (K) is untrue? The most straightforward way of explaining this is to endorse the following claim, which closely resembles (BCK_j):

> (BKC) We are all aware of the fact that, if an instance of (B) is true, and the corresponding instance of (K) is true, then the corresponding instance of (C) is true.

If (BKC) is true, then we are presumably aware of the fact that, if an instance of (B) is true, and the corresponding instance of (C) is untrue, then the corresponding instance of (K) is also untrue. And, if we are aware of this fact, then it is not surprising that, when we judge that an instance of (B) is true, and that the corresponding instance of (C) is untrue, we tend also to judge that the corresponding instance of (K) is untrue.

The problem with this explanation is that there are good reasons for thinking that (BKC) is false. One good way of seeing this is to consider the relationship between (BKC) and the following plausible principle about knowledge:

> (PK) If S can know, on the basis of the fact that p, that q, and the proposition that q obviously entails that r, then S can know, on the basis of the fact that p, that r.

When (BKC) is conjoined with this principle, it entails that many of our everyday claims about knowledge are false. One good way of seeing this is to focus on some of the claims about knowledge that were made in sections 3.5–3.7. In the situation described in section 3.6, it seemed clear that the following claim was true:

(K_{zp}) I can know, on the basis of the contents fact, that the animals in the enclosure are zebras.

But, if (BKC) and (PK) are both true, then (K_{zp}) is false. To see this, note first of all that, if (PK) is true, then (K_{zp}) implies the following claim:

(K_{mp}) I can know, on the basis of the contents fact, that the animals in the enclosure are not cleverly disguised mules.

Next, note that, if (BKC) is true, then (K_{mp}) implies the following claim:

(C_m) If the animals in the enclosure were cleverly disguised mules, then it would not be the case that they look to me like zebras, and are referred to as zebras by the sign in the enclosure.

Finally: note that, in the situation described in section 3.6, (C_m) is clearly false. From this, it follows that, if (BKC) and (PK) are both true, then, in the situation described in section 3.6, (K_{zp}) is false. But, it seems clear that, in the situation described in section 3.6, (K_{zp}) is *not* false. And, for this reason, it seems clear that, if we endorse (PK), then we should reject (BKC).

The most natural way of responding to this attack on (BKC) is to try to show that (PK) is false. But, it seems unlikely that a response of this kind will succeed. For, by responding to the attack in this way, we commit ourselves to endorsing 'abominable conjunctions' like the following:

(AC_K) I can know, on the basis of the contents fact, that the animals in the enclosure are zebras, but I cannot know, on the basis of this fact, that they are not cleverly disguised mules.

And, most of us are very strongly inclined to refrain from endorsing conjunctions of this kind.

Consequently, it is natural to say that (BKC) is false. But, if we say this, then we need another explanation for the truth of (BCK_j). One natural way of supplying this explanation is to appeal to the following contextualist claim about knowledge:[14]

> (CCK) When it is asserted (or implied) that an instance of (B) is true and that the corresponding instance of (K) is true, the standards for knowledge tend to be raised in such a way as to ensure that, if those instances of (B) and (K) are true then the corresponding instance of (C) is true.

This claim entails that, if an instance of (C) is untrue, and we assert that corresponding instances of (B) and (K) are both true, then the standards for knowledge tend to be raised in such a way as to make our assertion false. It also entails that, if an instance of (C) is untrue, and we assert that the corresponding instances of (B) and (K) are *not* both true, then the standards for knowledge tend to be raised in such a way as to make our assertion *true*. If the standards for knowledge vary in this way, then it is not surprising that, when we judge that an instance of (B) is true, and that the corresponding instance of (C) is untrue, we tend also to judge that the corresponding instance of (K) is untrue. So, if (CCK) is true, then it provides us with an explanation for the truth of (BCK_j).

Does this contextualist explanation suffer from the same problems as the explanation that appealed to (BKC)? On reflection, it is clear that it does not. For, unlike the (BKC) explanation, the contextualist explanation does not commit us to saying that (K_{zp}) is false. All that it commits us to saying is that, when it is asserted (or implied) that (K_{mp}) is true (or not true), the standards for knowledge tend to be raised in such a way as to make (K_{zp}) false. And, this claim is fairly plausible. In the situation

[14] DeRose 1995 defends a contextualist principle that is very similar to (CCK). In section 4 of chapter 4, we will say something about this form of contextualism. Some critics of DeRose have suggested that his contextualist principle should be replaced with something like (CCK). See for instance Williamson 2000, section 7.5.

described in section 3.6, it seems right to say that, once (K_{mp}) (or its negation) has been asserted (or implied) we can no longer truly assert (K_{zp}). What seems wrong is the claim that we can *never* truly assert (K_{zp}). The (BKC)-explanation commits us to making this claim, but the contextualist explanation does not. Consequently, we should opt for the contextualist explanation, by endorsing (CCK) rather than (BKC). By opting for this explanation, we generate another reason for believing (BEK).

5. Arguing for (KAB)

The upshot of the last two sections is that we have good reason to believe the following claim:

> (BEK) We are all aware of the fact that, if an instance of (B) is true, and the corresponding instance of (E) is untrue, then the corresponding instance of (K) also untrue.

In this section, I will argue that we also have good reason to believe this claim:

> (KAB) If S actively believes P, and S is aware of the fact that she does not know P, then S has a conclusive epistemic reason to stop actively believing P.

My argument for (KAB) will be divided into three stages. The first stage (sections 5.1 and 5.2) will describe three puzzles about knowledge and assertion, and will defend a solution to these puzzles that is due to Peter Unger. The second stage (section 5.3) will argue that, if we endorse Unger's solution, then we should also endorse a certain claim about knowledge and assertion; and, the third stage (section 5.4) will argue that, if we endorse this claim about knowledge and assertion, then we should also endorse (KAB).

5.1 *Three puzzles about knowledge and assertion*

5.1.1 *The assertions of sceptics*

When advocates of global scepticism assert that *there is absolutely nothing that we know*, or that *we cannot be sure of anything at all*, their assertions seem, somehow, to be inconsistent. But, when we look closely at these assertions, it seems perfectly clear that the assertions could have been true. There are possible worlds in which there is absolutely nothing that we know, and in which there is nothing of which we can be sure. So, why does it seem as if there are no such worlds? Why do the assertions of sceptics sound inconsistent?[15]

5.1.2 *Moore's paradox*

Consider the following sentences, which are often used to generate Moore's paradox:[16]

> (M_1) It is raining, but I do not believe that it is raining.
> (M_2) It is raining, but I do not know that it is raining.
> (M_3) It is raining, but I am not certain that it is raining.

When sentences of this kind are used to make assertions, the assertions sound inconsistent. But, when we look closely at the sentences, it seems perfectly clear that they could have been true. There are possible worlds in which it is raining, and in which I do not believe, do not know, and am not certain that it is raining. So, why does it seem as though there are no such worlds? Why do assertions of sentences like (M_1), (M_2) and (M_3) sound inconsistent?[17]

---

[15] Cf. Unger 1975:250-252.

[16] Normally, the paradox is generated by appeal to sentences about belief; but a number of authors – e.g. Moore 1962:277, Unger 1975:256-60 and Sorensen 1988:15-56 – have noted that similar paradoxes be generated by appeal to sentences about knowledge and certainty.

5.1.3 Challenges to assertion

When a subject, S, asserts that p, it often seems appropriate to attack their assertion by arguing that they do not *know* that p, or that they cannot be *certain* that p. This seems appropriate in spite of the fact that S has said *nothing at all* about knowing, or being certain that p. Generally speaking, it is not appropriate to attack an assertion by attacking some proposition other than the one that was asserted. So, why does it seem appropriate to attack S's assertion that p by attacking the proposition that S knows, or can be certain that p? Why can't S respond to such attacks by just pointing out that he only asserted *that p*, and did not say anything about *knowing*, or *being certain* of this proposition?[18]

5.2 Unger's solution

One very natural way of solving the above puzzles is to claim that, when we assert that p, we invariably imply that we *know* that p. The most prominent advocate of this solution is Peter Unger.[19] According to Unger, the apparent inconsistency of the assertion that *there is absolutely nothing that we know* is to be understood in terms of the actual inconsistency of the proposition that *I know that there is absolutely nothing that we know*. Similarly: the apparent inconsistency of the sentences (M_1), (M_2) and (M_3), above, is to be understood in terms of the actual inconsistency of the following sentences:

(M_{1k}) I know that it is raining, but I do not believe that it is raining.
(M_{2k}) I know that it is raining, but I do not know that it is raining.
(M_{3k}) I know that it is raining, but I am not certain that it is raining.

And, the fact that it is appropriate to attack S's assertion that p by attacking the proposition that *S knows that p* is to be understood in terms of the more general fact that an assertion can be attacked by attacking the propositions that the asserter implies, by making that assertion.

---

[17] Cf Unger 1975:256-60.
[18] Cf. Unger 1975:260-65

Unger's solution is very attractive, and seems to extend very naturally to a number of related puzzles.[20] Because of this, the solution has been widely endorsed.[21] In the next section, I will argue that the solution commits us to endorsing the following claim about knowledge and assertion:

(KA) If S is aware of the fact that she does not know that p, then S has a conclusive epistemic reason to refrain from asserting that p.

And, in section 5.4, I will argue that, by endorsing (KA), we commit ourselves to endorsing (KAB).

5.3 Defending (KA)

It is very plausible to suppose that, if it is not epistemically reasonable to *believe* that a certain proposition, P, is true, then it also is not epistemically reasonable to *assert*, or to *imply* that P is true. Because of this, it is natural to endorse the following claim about conclusive epistemic reasons:

(CER_1) If the fact that q is a conclusive epistemic reason for S to refrain from believing that p, then the fact that q is a conclusive epistemic reason for S to refrain from asserting, or implying that p.

It is fairly clear that, if it is not the case that p, then the fact that not-p is a conclusive epistemic reason for any subject, S, to refrain from believing that p. For this reason, (CER_1) commits us to endorsing the following claim:

(CER_2) If it is not the case that p, then the fact that it is not the case that p is a conclusive epistemic reason for any subject, S, to

[19] See Unger 1975, chp VI.
[20] For descriptions of some of these puzzles, see Unger 1975:260-65. See also DeRose 1991, where Unger's solution is applied to some puzzles about epistemic possibility.
[21] See e.g. Slote 1979, DeRose 1991 and Williamson 1996. Solutions to Moore's paradox which

refrain from asserting, or implying that p.

It is also clear that, if one has a conclusive epistemic reason to refrain from implying a proposition, and one would imply that proposition by asserting that p, then one has a conclusive epistemic reason to refrain from asserting that p. For this reason, (CER_2) commits us to endorsing the following claim:

(CER_3) If S would imply that q by asserting that p, and it is not the case that q, then the fact that it is not the case that q is a conclusive epistemic reason for S to refrain from asserting that p.

When this claim is conjoined with Unger's plausible claim that, by asserting that p, one invariably implies that one knows that p, it entails the following claim about knowledge and assertion:

(KA) If S is aware of the fact that she does not know that p, then S has a conclusive epistemic reason to refrain from asserting that p.

This claim about knowledge and assertion is obviously very similar to the principle (KAB) which we are trying to defend. In order to move from (KA) to (KAB), we should presumably appeal to something like the following claim about reasons for assertion and belief:

(RAB) If S has a conclusive epistemic reason to refrain from asserting that p, then S has a conclusive epistemic reason to refrain from actively believing that p.

It is very plausible to suppose that something like (RAB) is true. But, as it stands, (RAB) is open to a certain objection. In section 5.4, we will argue that, by modifying (RAB), we can save it from this objection. The modified principle that we will defend is weaker than (RAB), but is still strong enough to take us from (KA) to

resemble Unger's are endorsed in Moore 1962:277 and Sorensen 1988:15-56.

(KAB).

5.4 Modifying (RAB)

One good way of seeing that (RAB) is in need of modification is to focus on cases in which people conversationally imply things that they know to be false.[22] Suppose that Professor X has just been asked to write a letter of recommendation for student Y, who is applying for an academic job. Although X knows that Y is an intelligent student, he wants to imply that this is not the case; so, in his letter, X asserts, pointedly, that Y is not the most intelligent student he has ever taught. If the circumstances here are normal, then, by asserting this, X will imply that Y is not an intelligent student. So, since Y *is* an intelligent student, and X is aware of this, the following principle, which was defended in section 5.3: commits us to saying that, in these circumstances, X has a conclusive epistemic reason to refrain from asserting that Y is not the most intelligent student he has ever taught:

> (CER_3) If S would imply that q by asserting that p, and it is not the case that q, then the fact that it is not the case that q is a conclusive epistemic reason for S to refrain from asserting that p.

However: it seems clear that, in these circumstances, X does not – or at least, need not – have a conclusive epistemic reason to refrain from *believing* that Y is not the most intelligent student he has ever taught. And consequently, it seems clear that, in its current form, (RAB) is false.

The way to save (RAB) from this kind of counterexample is to focus on the fact that there are circumstances in which Professor X's assertion would *not* imply that Y is not an intelligent student. Because there are such circumstances, we can save (RAB) by rewriting it in the following way:

---

[22] I am here using the term 'conversationally imply' in the technical sense that was introduced in Grice 1975.

(RAB_2) If S is aware of some fact which would in *any* circumstances be a conclusive epistemic reason for her to refrain from asserting that p, then S has a conclusive epistemic reason to refrain from actively believing that p.

When (RAB) is rewritten in this way, it still licenses the move from (KA) to (KAB). For, (KA) can be read as saying that *any* circumstances in which S is aware of the fact that she does not know that p are circumstances in which she has a conclusive reason to refrain from asserting that p. If knowledge was a cancellable implication of assertion, then it would not be legitimate to read (KA) in this way. But, as Slote has noted, knowledge is *not* a cancellable implication of assertion.[23] If it was, then there would be some circumstances in which Moorean assertions like *p but I do not know that p* lose their apparent inconsistency. But, on reflection, it seems clear that there are no circumstances of this kind.

5.5 Summary

The upshot of the last four subsections is that we have good reason to believe both of the following claims:

(KA) If S is aware of the fact that she does not know that p, then S has a conclusive epistemic reason to refrain from asserting that p.

(RAB_2) If S is aware of some fact which would in *any* circumstances be a conclusive epistemic reason for her to refrain from asserting that p, then S has a conclusive epistemic reason to refrain from actively believing that p.

When these two claims are conjoined, they entail the following claim:

(KAB) If S actively believes P, and S is aware of the fact that she does not know P, then S has a conclusive epistemic reason to stop actively believing P.

So, we can now conclude that there is good reason to believe this claim. In sections 3 and 4, we showed that there is also good reason to believe the following claim (where (B), (E) and (K) refer to the schematic sentences described in section 8):

(BEK) We are all aware of the fact that, if an instance of (B) is true, and the corresponding instance of (E) is untrue, then the corresponding instance of (K) is also untrue.

And, in section 2, we showed that acceptance of (BEK) and (KAB) commits us to endorsing the LOK explanation. So, we can now conclude that the LOK explanation ought to be endorsed. In section 6, we will use the LOK explanation to generate a problem for the anti-sceptical objection that we described in chapter one.

6. The LOK explanation and the anti-sceptical objection

Before we can show that the LOK explanation generates a problem for the anti-sceptical objection, we must trace out some of the consequences of this explanation. Sections 6.1 and 6.2 will argue that, if the LOK explanation is correct, then there is a form of evidentialism that discredits many beliefs, and section 6.3 will argue that, by appealing to this form of evidentialism, we can generate a problem for the anti-sceptical objection.

6.1 Generalising the LOK explanation

There are many everyday cases in which someone challenges a certain belief by asking the believer *why* he or she has that belief. In section 1 of this chapter, we described three cases of this kind. Each of these cases conforms to the following pattern (which we will call *the alpha pattern*):

(i) First, person A tells person B that p.

[23] Slote 1979:179.

(ii) Then, B asks A why he believes that p.

(iii) A responds by saying, sincerely, that he believes that p because q.

(iv) B then makes it clear to A that the fact that q does not *show*, or *establish* that p.

And, in each of these cases, it seems right to say that B's questions and comments give A a conclusive epistemic reason to stop actively believing that p.

It is easy to imagine other cases that conform to the alpha pattern. And, it is interesting to note that, in virtually *all* of these cases, B's questions and comments seem to provide A with a conclusive epistemic reason to stop actively believing that p. This fact about alpha-pattern cases stands in need of explanation. If the LOK explanation is correct, then we should surely explain it in the following way:

In virtually every case that conforms to the alpha-pattern, B's questions and comments make A aware of the fact that the following claims are both true with respect to him:

(B_p) (I believe that p) because q.

($\neg E_p$) The fact that q does not establish that p.

And, in virtually every case of this kind, A is already aware of the fact that, if (B_p) and ($\neg E_p$) are true with respect to him, then the following claim is also true with respect to him:

($\neg K_p$) I do not know that p.

So, in virtually every alpha-pattern case, B makes A aware of the fact that he does not know that p. Because of this, virtually every alpha-pattern case is a case in which B gives A a conclusive epistemic reason to stop actively believing that p.

In what follows, we will refer to this as *the generalised LOK explanation*. The arguments of sections 2–5 make it clear that this explanation is one that we ought to endorse. In section 6.2, we will argue that, if the generalised LOK explanation is true, then there is a form of evidentialism that discredits many beliefs. Section 6.3 will use this form of evidentialism to generate a problem for the anti-sceptical objection that we discussed in chapter one.

6.2 *Explanatory evidentialism*

In section 6.2.1, I will argue that, if the generalised LOK explanation is true, then we are all in a position to know the following doctrine:

(E_E) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not explained by establishing evidence, then S's belief in P does not constitute knowledge.

Section 6.2.2 will argue that this doctrine is a form of evidentialism, and section 6.2.3 will argue that it discredits many beliefs.

6.2.1 *(E_E) and the generalised LOK explanation*

If the generalised LOK explanation is true, then the following claim about alpha-pattern cases is also true:

(APC) In virtually every alpha-pattern case, A is aware of the fact that, if (B_p) and ($\neg E_p$) are both true with respect to him, then ($\neg K_p$) is also true with respect to him.

The fact about alpha-pattern cases itself stands in need of explanation. For, there are many different kinds of alpha-pattern case, involving subjects with very different kinds of background knowledge. How is it that, in virtually *all* of these cases, person A is aware of the truth of the conditional to which (APC) refers?

One natural way of explaining this fact is to endorse the following claim:

(APC_2) In virtually every alpha-pattern case, it is *a conceptual truth* that, if (B_p) and ($\neg E_p$) are both true with respect to A, then ($\neg K_p$) is also true with respect to A.

It is hard to see how one could explain the truth of (APC) without endorsing something like this claim. Because of this, we have good reason to believe (APC_2). But, the truth of (APC_2) is itself something that stands in need of explanation.

The most natural way of explaining the truth of (APC_2) is to say that there is a broad conceptual truth that entails all of the conceptual truths to which (APC_2) refers. One simple claim which entails all of the conceptual truths to which (APC_2) refers is the following (where S's belief that p is *explained by establishing evidence* iff there is some q such that (i) (S believes that p) because q, and (ii) the fact that q establishes that p):

(E) For any subject, S, and proposition P: if S believes P, and S's belief in P is not explained by establishing evidence, then S does not know P.

There are a few propositions for which (E) may not hold. These propositions are all simple necessary truths, like the proposition that *everything is self-identical*, and the proposition that $2 + 2 = 4$. Although it is clear that we know these propositions, it is by no means clear that we believe them in the way that (E) requires. Because of this, it seems best to explain the truth of (APC_2) by appealing to the following qualified version of (E):

(E_E) With few exceptions: every proposition P, is such that, for any subject, S: if S believes P, and S's belief in P is not explained by establishing evidence, then S does not know P.

It is hard to see how we could explain the truth of (APC_2) without saying that (E_E) is a conceptual truth. Because of this, we have good reason to believe that (E_E) is a conceptual truth. But, if (E_E) is a conceptual truth, then, surely, we are all in a position to *know* (E_E). In the next subsection, I will argue that (E_E) is a form of evidentialism.

6.2.2 *From (E_1) to (E_E)*

A doctrine, D, is a *form of evidentialism* iff D is either identical to the following doctrine, or can be generated by making an admissible amendment to this doctrine:

(E_1) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P is epistemically irrational.

In what follows, we will show that there is an admissible amendment to (E_1) which generates (E_E). By showing this, we will show that (E_E) is a form of evidentialism.

One way of amending (E_1) is to amend the doxastic defect to which it refers. Amendments of this kind are admissible iff the defect to which they appeal is suitable – where a doxastic defect is said to be *suitable* iff it satisfies the following constraint:

(SDD) For any subject, S, and proposition P: if S believes P, and S's belief in P has defect D, then the fact that this belief has D is a conclusive epistemic reason for S to stop believing P.

The arguments of section 5 show that the defect of *not constituting knowledge* satisfies (SDD). So, we can now conclude that the following doctrine is a form of evidentialism:

(E_2) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P *does not constitute knowledge*.

Another way of amending (E_1) is to amend the evidential requirement that it imposes on our beliefs. One admissible amendment of this kind is the amendment which generates the following doctrine:

(E_1') With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S 's belief in P is *not based on sufficient evidence*, then S's belief in P is epistemically irrational.

Any amendment to (E_1) which can be brought about by making a series of admissible amendments to this doctrine is itself admissible. So, by combining the amendment which generates (E_2) with the amendment that generates (E_1'), we can generate the following form of evidentialism:

(E_3) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not based on sufficient evidence, then S's belief in P does not constitute knowledge.

A third admissible way of amending (E_1) is to replace its constituent terms with terms that can reasonably be believed to be their synonyms. Any amendment to (E_1) which can be brought about by making a series of admissible amendments to this doctrine is itself admissible; so, any amendment to (E_1) which can be brought about by first of all replacing (E_1) with (E_1'), and then replacing terms in (E_1') with terms that can reasonably be believed to be their synonyms, is an admissible amendment to (E_1). It is reasonable to believe that the following schematic sentence:

($\neg$BE) S's belief in P is not based on sufficient evidence.

is synonymous with this sentence:

($\neg$EEE) S's belief in P is not explained by establishing evidence.

Because it is reasonable to believe this, it is possible to generate the following doctrine by making an admissible amendment to (E_1):

(E_E) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not explained by establishing evidence, then S's belief in P does not constitute knowledge.

Consequently, we can conclude that (E_E) is a form of evidentialism. In what follows, we will refer to it as *explanatory evidentialism.*

6.2.3 Explanatory evidentialism discredits many beliefs

In the last two subsections, we have argued that we are all in a position to know a certain form of evidentialism, which we are calling explanatory evidentialism. In this section, we will use the results of these sections to argue that explanatory evidentialism discredits many defective beliefs. We will do this by focusing on the job interview case that we described in section 1. In what follows, we will argue that explanatory evidentialism discredits John's job belief, and will thereby argue that there are many defective beliefs that are discredited by this form of evidentialism.

When I say that a form of evidentialism, E, *discredits* S's belief that p, what I mean is that the following claims are all true:

(D1') S is in a position to know E, and is in a position to know that her belief that p is not an exception to it.

(D2') If S is not aware of the fact that her belief that p does not satisfy the evidential requirement that E imposes, then we can make S aware of this fact.

(D3') (D1') and (D2') are true regardless of whether S's belief that p is true, and regardless of S's belief that *she knows that p* can be rebutted in a non-evidential way. They are also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

The arguments of section 6.2.1 show that we are all in a position to know explanatory evidentialism; and, it seems clear that John is in a position to know that his job belief is not an exception to this form of evidentialism. Consequently, it seems clear that, in the job interview case, the following claim is true:

(D1$_j$') John is in a position to know explanatory evidentialism, and is in a position to know that his job belief is not an exception to it.

It also seems clear that, in the job interview case, the following pair of claims are true:

(D2$_j$') If John is not aware of the fact that his job belief does not satisfy the evidential requirement that is imposed by explanatory evidentialism, then we can make John aware of this fact

(D3$_j$') (D1$_j$') and (D2$_j$') are true regardless of whether John's job belief is true, and regardless of whether John's belief that *his job belief constitutes knowledge* can be rebutted in a non-evidential way. They are also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

Consequently, it seems clear that, in the job interview case, explanatory evidentialism discredits John's job belief. And, since there are many actual cases that closely resemble the job interview case, it seems clear that there are many other beliefs that are discredited by this form of evidentialism. By appealing to these facts about explanatory evidentialism, we can generate a problem for the anti-sceptical objection that we described in chapter one. In section 6.3, I will describe this problem. Chapter three will then describe a way in which the problem can be resolved.

6.3 A problem for the anti-sceptical objection

There are two stages of the anti-sceptical objection. The first stage is an argument for the following claim:

(1) EBG is true only if there is a form of evidentialism that discredits all deity beliefs.

and the second stage is an argument for this claim:

(2) There is no form of evidentialism that discredits all deity beliefs.

To see that the results of section 6.2 generate a problem for the anti-sceptical objection, we need to focus on the second stage of this objection. The most important part of this stage of the objection is an argument for the claim that the following form of evidentialism leads to scepticism:

(E_I) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P is epistemically irrational.

According to the anti-sceptical objector, (E_I) leads to scepticism because we do not have evidence for the negations of *sceptical hypotheses*, like the hypothesis that *I am a brain in a vat*. If the objector's argument succeeds, then it shows that most forms of evidentialism lead to scepticism, and thus do not discredit all deity beliefs. The only forms of evidentialism that seem likely to escape the objector's argument are those which impose a fairly *weak* evidential requirement on our beliefs. And, it is difficult to argue that deity beliefs fail to satisfy these weak requirements without appealing to arguments for the falsity of certain metaphysical deity beliefs, or to non-evidential attacks on certain epistemic deity beliefs: so, even these forms of evidentialism seem unlikely to discredit all deity beliefs.

One form of evidentialism which seems vulnerable to the second stage of the anti-sceptical objection is the following doctrine, which we are calling *explanatory evidentialism*:

(E_E) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S's belief in P is not explained by establishing evidence, then S's belief in P does not constitute knowledge.

This form of evidentialism seems vulnerable to the objection because the evidential requirement that it imposes seems not to be satisfied by our beliefs in the negations of sceptical hypotheses. My belief that *I am not a brain in a vat* (henceforth: that I am not a *BIV*) seems to be explained by the fact that *I believe lots of propositions that entail that I am not a BIV.* And, it seems clear that this fact does not establish that I am not a BIV. For, if I was a BIV, then I would believe all of the things that I actually believe. And, this seems to show that *no* fact about the things that I believe can establish that I am not a BIV.

To see that explanatory evidentialism is not in fact vulnerable to the anti-sceptical objection, we need to recap on some of the things that we said about establishment in section 4.1. In this section, we discussed the fact that instances of the following claim:

(E) The fact that q establishes that p.

seem always to imply corresponding instances of this claim:

(C) If it hadn't been the case that p, then it would not have been the case that q.

The upshot of this section was that this implication does not always hold. In particular: we concluded that the implication does not hold in cases where the proposition that p is the negation of a *sceptical hypothesis*. So, for instance, we concluded that, in ordinary contexts, the following instance of (E) is true:

(E_m) The fact that *the animals in the zebra enclosure look like zebras* establishes that *these animals are not cleverly disguised mules.*

in spite of the fact that, in such contexts, the corresponding instance of (C) seems false. Similarly, we concluded that, in ordinary contexts, the following instance of

(E) is true:

> (E_s) The fact that *Jim is wearing a suit that looks new* establishes that *Jim has just* bought, *rather than stolen, a new suit.*

in spite of the fact that, in such contexts, the corresponding instance of (C) seems false. If the argument of section 4.1 succeeds, then it allows us to say that, in ordinary contexts, the following instance of (E) is true:

> (E_b) The fact that *I believe lots of propositions that entail that I am not a BIV* establishes that *I am not a BIV.*

even though it seems clear that, in such contexts, the corresponding instance of (C) is false. Because of this, it allows us to say that, in ordinary contexts, my belief in the negation of the BIV hypothesis *does* satisfy the evidential requirement that explanatory evidentialism imposes. And, consequently, it provides us with a way of responding to the anti-sceptical objector's argument for the claim that (E_E) leads to scepticism.

Advocates of the anti-sceptical objection may respond to this problem by changing tack. Instead of trying to argue that explanatory evidentialism leads to scepticism, they may try to argue that the evidential requirement that it imposes is too weak. If my belief in the negation of the BIV hypothesis satisfies the evidential requirement that explanatory evidentialism imposes, then this requirement is, at least sometimes, much easier to satisfy than it seems. By appealing to this fact, the anti-sceptical objector may try to show that the following pair of claims do not hold with respect to every subject, S, who has deity beliefs:

> ($D2_{EE}$) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that explanatory evidentialism imposes, then we can make S aware of this fact.

($D3_{EE}$) ($D2_{EE}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way.

If the objector can show this, then she will show that explanatory evidentialism is too weak to discredit all deity beliefs.

The problem with this strategy is that there are many beliefs that are discredited by explanatory evidentialism in spite of the fact that the evidential requirement imposed by this form of evidentialism is sometimes much easier to satisfy than it seems. One belief of this kind is John's job belief. To show that explanatory evidentialism is too weak to discredit all deity beliefs, the anti-sceptical objector must appeal to features of deity beliefs which are not shared by John's job belief. At present, the anti-sceptical objection does not appeal to features of this kind – so, at present, this objection does not show that explanatory evidentialism is too weak to discredit all deity beliefs.

The upshot of this chapter is that, in its current form, the anti-sceptical objection does not show that explanatory evidentialism discredits all deity beliefs. If this problem with the objection cannot be rectified, then the objection will fail. In chapter three, I will rectify the problem by constructing a new version of the anti-sceptical objection. Chapter four will then show that this new objection is itself vulnerable to a certain problem, which will be rectified in chapter five.

Chapter Three

The second anti-sceptical objection

In chapter two, I argued that advocates of the anti-sceptical objection (ASO) face a serious problem. In this chapter, I will construct a new version of the ASO, which is invulnerable to this problem. Before I can start to construct this new objection, I need to argue for the following claim:

(M) If there is some subject, S, whose metaphysical deity beliefs imply that some of his deity beliefs are explained by sensitive evidence, then explanatory evidentialism does not discredit all deity beliefs.

Section 1 clarifies this claim, sections 2 and 3 argue for it, and section 4 uses it to construct an improved version of the ASO, which I call *the second ASO*.

1. Clarifying (M)

Before we can defend (M), we must clarify the terminology that it employs. In particular: we need to say something about what it is for a belief to *explained by sensitive evidence*. When I say that S's belief that p is explained by sensitive evidence, I mean that an instance of the following claim is true:

(ESE) (S believes that p) because q, and if it hadn't been the case that p, then it wouldn't have been the case that q.

Because of this, we can clarify (M) by rewriting it in the following way:

(M') If there is some subject, S, whose metaphysical deity beliefs imply that there is some sentence, 'd', which states the content of one of S's deity beliefs, and for which an instance of the following claim is true:

(ESE) (S believes that d) because q, and if it hadn't been the case that d, then it wouldn't have been the case that q.

then explanatory evidentialism does not discredit all deity beliefs.

In the next two sections, I will give a two-step argument for (M'). Section 2 will argue that, if the antecedent of (M') is true, then, under ordinary standards for establishment, the following claim is also true:

(M_{ae}') There is some subject, S, whose metaphysical deity beliefs imply that there is some sentence, 'd', which states the content of one of S's deity beliefs, and for which an instance of the following claim is true:

(EEE) (S believes that d) because q, and the fact that q establishes that d.

And, section 3 will argue that, if (M_{ae}') is true under ordinary standards for establishment, then the consequent of (M') is true.

2. From (M') to (M_{ae}')

Chapter 2 argued that, when people judge that an instance of the following claim is true:

(E) The fact that p establishes that q

they tend also to judge that the corresponding instance of this claim is true:

(C) If it had not been the case that q, then it would not have been the case that p.[1]

It is interesting to note that the reverse tendency also seems to hold. When people judge that an instance of (C) is true, they tend also to judge that the corresponding instance of (E) is true. One good way of seeing this is to appeal to examples.

Last night, it snowed, and this morning, there is a layer of snow on my garden. Earlier this morning, a bird landed on the garden, and left some footprints in the snow. It seems clear that the following instance of (C) is true:

(C_b) If a bird hadn't landed on my garden, then there would not have been bird footprints in the snow.

And, it seems equally clear that this instance of (E) is true:

(E_b) The fact that *there are bird footprints in the snow* establishes that a bird landed on my garden.

If there was a trickster in the area who was in the habit of leaving fake bird footprints on people's gardens, then it would be less clear that (E_b) is true. But, it would also be less clear that (C_b) is true, which supports the hypothesis that the plausibility of (E_b) depends on the plausibility of (C_b).

Last night, I left a saucer of milk in the kitchen. During the night, my cat came into the house and drank the milk. It seems clear that the following instance of (C) is true:

(C_c) If the cat hadn't drunk the milk, then the saucer wouldn't have been empty this morning.

And, it seems equally clear that this instance of (E) is true:

(E_c) The fact that *the saucer was empty this morning* establishes that the cat drank the milk.

[1] See sections 3.1-3.3 of chapter two.

If the kitchen was accessible to other animals – e.g. to hedgehogs – then it would be less clear that (E_c) is true. But, it would also be less clear that (C_c) was true, which supports the hypothesis that the plausibility of (E_c) depends upon that of (C_c).

This morning, I parked my car in a multi-storey car park. Later in the day, the car was stolen; so, when I returned, the car was not in the car park. Since no-one but me has keys to the car, the following instance of (C) is true:

(C_p) If the car hadn't been stolen, then the car would have been in the car park when I returned.

The corresponding instance of (E) also seems to be true:

(E_p) The fact that *the car was not in the car park when I returned* establishes that the car had been stolen.

If someone else – e.g. my father – had keys to the car, then it would be less clear that (E_p) is true. But, it would also be less clear that (C_p) is true, which supports that hypothesis that (E_p) owes its plausibility to (E_p).

It is natural to conclude, on the basis of the above examples, that, when people judge that an instance of (C) is true, they tend also to judge that the corresponding instance of (E) is true. It is also natural to conclude, on the basis of these examples, that instances of (C) *obviously imply* the corresponding instances of (E). But, there are good reasons to refrain from endorsing this second conclusion. One good way of appreciating these reasons is to return to the following instances of (C) and (E), which we discussed above:

(C_c) If the cat hadn't drunk the milk, then the saucer wouldn't have been empty this morning.

(E_c) The fact that *the saucer was empty this morning* establishes that the cat drank the milk.

It is natural to say that (C_c) obviously implies (E_c). But, there is also reason to think that this implication does not hold. To see this, note that (C_c) seems *not* to imply the following claim:

(E_c*) The fact that *the saucer was empty this morning* establishes that the milk was not drunk by a hedgehog that somehow managed to sneak into the house.

Although this claim is not implied by (C_c), it *does* seem to be implied by (E_c). For, the proposition that *the cat drank the milk* entails that *the milk was not drunk by a hedgehog that somehow managed to sneak into the house*; and, establishment seems to be closed under entailment. If (E_c*) is implied by (E_c), but not (C_c), then (E_c) is not implied by (C_c).[2] Because of this, there are problems for those who wish to say that instances of (C) obviously imply corresponding instances of (E).

The best way of resolving these problems is to qualify the claim that instances of (C) obviously imply corresponding instances of (E). The contextualist arguments of chapter two[3] suggest that we should qualify this claim in the following way:

(CE_c) Under ordinary standards for establishment, instances of (C) obviously imply corresponding instances of (E).

If we do qualify it in this way, then we can respond to the above problems by saying that, under ordinary standards for establishment, (C_c) *does* imply (E_c*). The apparent failure of this implication can be explained by saying that, when (E_c*) is mentioned, the standards for establishment are raised.

Since we have already accepted, in chapter two, that the standards for establishment vary with context, we have good reason to accept (CE_c). From (CE_c), it follows that, under ordinary standards for establishment, the following claim:

---

[2] Here, I assume that the relevant kind of implication is transitive. We can stipulate that this is so by stipulating that the relevant kind of implication is material implication. This stipulation does not make the claims of this section any less plausible.

[3] These arguments are given in section 4.1 of chapter two.

(M_a') There is some subject, S, whose metaphysical deity beliefs imply that there is some sentence, 'd', which states the content of one of S's deity beliefs, and for which an instance of the following claim is true:

(ESE) (S believes that d) because q, and if it hadn't been the case that d, then it wouldn't have been the case that q.

implies this claim:

(M_{ae}') There is some subject, S, whose metaphysical deity beliefs imply that there is some sentence, 'd', which states the content of one of S's deity beliefs, and for which an instance of the following claim is true:

(EEE) (S believes that d) because q, and the fact that q establishes that d.

In section 3, I will show that, if (M_{ae}') is true under ordinary standards for establishment, then the following claim is also true:

(M_c') Explanatory evidentialism does not discredit all deity beliefs.

By showing this, I will show that the following claim is true (since (M_a') and (M_c') are the antecedent and consequent of this claim):

(M') If there is some subject, S, whose metaphysical deity beliefs imply that there is some sentence, 'd', which states the content of one of S's deity beliefs, and for which an instance of the following claim is true:

(ESE) (S believes that d) because q, and if it hadn't been the case that d, then it wouldn't have been the case that q.

then explanatory evidentialism does not discredit all deity beliefs.

3. From (M_{ae}') to (M_c')

If, under ordinary standards for establishment, the following claim is true:

(M_{ae}') There is some subject, S, whose metaphysical deity beliefs imply that there is some sentence, 'd', which states the content of one of S's deity beliefs, and for which an instance of the following claim is true:

(EEE) (S believes that d) because q, and the fact that q establishes that d.

then this claim is also true:

(M_c') Explanatory evidentialism does not discredit all deity beliefs.

To see this, we need to recap, briefly, on the definition of discrediting that we gave in chapter one. According to this definition, a form of evidentialism, E, discredits all deity beliefs iff it is the case that, for every subject, S, who has deity beliefs, the following three claims are all true:

(D1) S is in a position to know E, and is in a position to know that her deity beliefs are not exceptions to it.

(D2) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that E imposes, then we can make S aware of this fact.

(D3) (D1) and (D2) are true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way. They are also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

From this definition, we can infer that, if explanatory evidentialism discredits all deity beliefs, then, for every subject, S, who has deity beliefs, the following pair of claims are true:

(C_1) If S is not aware of the fact that her deity beliefs are not explained by establishing evidence, then we can make S aware of this fact.

(C_2) (C_1) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether the standards for establishment have been raised.

If (C_1) and (C_2) are true for every subject, S, then, under ordinary standards for establishment, there is no subject, S, for whom the following claim is true:

(C_3) If S's metaphysical deity beliefs are all true, then some of S's deity beliefs are explained by establishing evidence.

But, if the following claim is true, under ordinary standards for establishment:

(M_{ae}) There is some subject, S, whose metaphysical deity beliefs imply that some of his deity beliefs are explained by establishing evidence.

then there *is* some subject, S, for whom (C_3) is true. Consequently, we can conclude that, if (M_{ae}) is true, under ordinary standards for establishment, then explanatory evidentialism does not discredit all deity beliefs. And, since (M_{ae}) can be rewritten in the following way:

(M_{ae}') There is some subject, S, whose metaphysical deity beliefs imply that there is some sentence, 'd', which states the content of one of S's deity beliefs, and for which an instance of the following claim is true:

(EEE) (S believes that d) because q, and the fact that q establishes that d.

we can conclude that, if (M_{ae}') is true under ordinary standards for establishment, then the following claim is true:

(M_c') Explanatory evidentialism does not discredit all deity beliefs.

When the conclusion of this section is conjoined with the conclusion of the last section, it entails that the following claim is true:

(M) If there is some subject, S, whose metaphysical deity beliefs imply that some of his deity beliefs are explained by sensitive evidence, then explanatory evidentialism does not discredit all deity beliefs.

In the next section, I will show that, by appealing to this claim, we can construct an improved version of the anti-sceptical objection.

4. The second anti-sceptical objection

In this section, I will describe an improved version of the anti-sceptical objection (henceforth: the ASO). To distinguish it from the objection that we described in chapter one, I will refer to it as the second ASO; or, for short, as ASO_2. The second ASO is a four-stage argument which aims to give advocates of EBG a conclusive epistemic reason to stop actively believing this claim. In what follows, I will describe this argument by describing the four arguments from which it is composed.

The first stage of ASO_2 is the argument that we gave in chapter one for the following claim:

(1) If EBG is true, then there is a form of evidentialism that discredits all deity beliefs.

Nothing that we said in chapter two threatens the argument that we gave for (1); so, we can safely retain this argument as a part of ASO_2.

The second stage of ASO_2 is an argument for a qualified version of this claim:

(2) There is no form of evidentialism that discredits all deity beliefs.

We can generate this argument by focusing on the second stage of the original ASO. According to this stage of the objection, virtually every form of evidentialism either leads to scepticism, or is too weak to discredit deity beliefs. In chapter two, we showed that explanatory evidentialism does not lead to scepticism, and may not be too weak to discredit all deity beliefs; but, it is hard to think of any other form of evidentialism which shares these features. Because of this, the second stage of the original ASO gives us reason to believe the following, qualified version of claim (2):

(2') If explanatory evidentialism does not discredit all deity beliefs, then there is no form of evidentialism that discredits all deity beliefs.

The second stage of ASO_2 is the argument that we have just given for this claim.

The third stage of ASO_2 is a three-part argument for the following claim:

(3') If there is some subject, S, whose metaphysical deity beliefs imply that some of his deity beliefs are explained by sensitive evidence, then EBG is not true.

The first part of this argument points out that claims (1) and (2') jointly entail the following claim:

(3a') If explanatory evidentialism does not discredit all deity beliefs, then EBG is not true.

The second part uses the arguments of sections 1-3 to defend this claim:

(M) If there is some subject, S, whose metaphysical deity beliefs imply that some of his deity beliefs are explained by sensitive evidence, then explanatory evidentialism does not discredit all deity beliefs.

And, the third part points out that (3a') and (M) jointly entail (3').

The fourth, and final stage of ASO_2 is an argument for this claim:

(4') Under ordinary standards for knowledge, EBG is not known to be true.

To state this argument, we need to focus on the following claim:

(D) There is a community of people who all believe in a certain deity (who we will call *D*). Many of the things that the members of this community believe about D are believed by them because they are written in a certain book (which we will call *the D-book*). According to this community of D-believers, the D-book is, in a certain sense, a book that was written by D himself. The D-believers also think that D would never write a book which said something untrue (and they think this for the sense of 'write' which they use when they say that *the D-book was written by D*).

The fourth stage of ASO_2 argues for (4') by arguing for the following pair of claims:

(4a') If (D) is true, then EBG is not true.

(4b') If (D) is not true, then, under ordinary standards for knowledge, EBG is not known to be true.

In the next five subsections, I will state the fourth stage of ASO_2, by stating the arguments from which it is composed.

4.1 Arguing for (4a')

If (D) is true, then there are many true instances of the following claim in which 'S' is substituted for the name of a D-believer, and 'd' is substituted for a sentence which states the content of one of S's beliefs about D:

(B_d) (S believes that d) because the D-book says that d.

All of these instances of (B_d) are such that the subject, S, to whom they refer, believes (i) that the D-book was, in a certain sense, written by D himself, and (ii) that D would never write a book which said something untrue. Because of this, all of them are such that the subject, S, to whom they refer, has metaphysical deity beliefs which imply that the following claim is true:

(C_d) If it hadn't been the case that d, then the D-book would not have said that d.

And, consequently, it seems clear that, if (D) is true, then there are subjects who make the antecedent of the following claim true:

(3') If there is some subject, S, whose metaphysical deity beliefs imply that some of his deity beliefs are explained by sensitive evidence, then EBG is not true.

The third stage of ASO_2 shows that claim (3') is true; so we can now conclude that, if (D) is true, then the consequent of (3') is also true; or, equivalently, that the following claim is true:

(4a') If (D) is true, then EBG is not true.

4.2 Arguing for (4b')

Before we can argue for (4b'), we need to argue for two other claims. The first of these claims can be stated in the following way:

(EP_1) If (D) is not true, then it could easily have been true.

And, the second can be stated as follows:

(EP_2) If it could easily have been the case that p, then, under ordinary standards for establishment, instances of the following claim:

($\neg$C) If it had not been the case that p, then it would still have been the case that q.

imply corresponding instances of this claim:

($\neg$E) The fact that q does not establish that p.

In the next two subsections, I will argue for these claims. Section 4.5 will then use these claims to argue for (4b').

4.3 Arguing for (EP_1)

There are certain actual communities which closely resemble the community described in claim (D). One such is the community of Christians. Many of the things that Christians believe about God are such that Christians believe them because (or at least: partly because) the Bible says that they are true. And, many Christians seem to think (a) that the Bible was, in a certain sense, written by God, and (b) that God would never write a book which said something untrue.

It may turn out that the community of Christians does not conform exactly to our description of the D-believers. And, it may turn out that there is no other community which conforms exactly to this description. But, it seems clear that there are a number of communities that come *close* to conforming to this description. And, for this reason, it is natural to endorse the following claim:

(EP_1) If (D) is not true, then it could easily have been true.

4.4 *Arguing for (EP_2)*

In chapter two, we saw that, under ordinary standards for establishment, instances of the following claim:

> ($\neg$C) If it hadn't been the case that p, then it would still have been the case that q.

do not always imply instances of this claim:

> ($\neg$E) The fact that q doesn't establish that p.[4]

In spite of this, it seems clear that instances of ($\neg$C) do *sometimes* imply instances of ($\neg$E), under ordinary standards for establishment. To see this, recall the case of the broken thermometer that we discussed in section 1.2 of chapter 2. In this case, it seems clear that the following claim:

> ($\neg C_t$) If Joe's temperature had not been normal, then the thermometer would still have said that it was normal.

implies this claim:

> ($\neg E_t$) The fact that *the thermometer says that Joe's temperature is normal* does not establish that Joe's is normal.

and that it does so under ordinary standards for establishment. Because of this, we should accept that instances of ($\neg$C) do sometimes imply instances of ($\neg$E), under such standards.

What is the difference between the cases in which instances of ($\neg$C) imply instances of ($\neg$E), under ordinary standards for establishment, and the cases in which they don't? One good way of isolating this difference is to return to the zebra example that we discussed in section 3.2 of chapter two. In this example, I am visiting my local zoo, and am looking into the zebra enclosure. Inside the enclosure, there are striped animals which look to me like zebras, and a sign which says that these animals *are* zebras. Under ordinary standards for establishment, this fact about the contents of the enclosure – which we called *the contents fact* – establishes that the animals in the enclosure are zebras. It also establishes that these animals are not *cleverly disguised mules*, since the claim that they are zebras entails that they are not cleverly disguised mules. Because it establishes that these animals are not cleverly disguised mules, we can conclude that, in this case, the following instance of (¬C):

> ($\neg C_m$) If the animals in the enclosure had been cleverly disguised mules, then they would still have looked to me like zebras, and would still have been referred to as zebras by the sign in their enclosure.

does not imply the corresponding instance of (¬E). What is the difference between this case and the thermometer case, in which an instance of (¬C) *does* imply the corresponding instance of (¬E)?

The most striking difference between these two cases is that in the first case, but not the second, the antecedent of the relevant instance of (¬C) *could easily have been true*. In the thermometer case, it could easily have been the case that the patient's temperature was not normal; but, in the zebra case, it could *not* easily have been the case that the animals in the enclosure were cleverly painted mules.[5] This fact about the two examples suggests that, under ordinary standards for establishment, instances of (¬C) imply corresponding instances of (¬E) iff the antecedent of (¬C) could easily have been true. If this suggestion is right, then the following claim is true:

> (EP_2) If it could easily have been the case that not-p, then, under ordinary standards for establishment, instances of the following claim:

---

[4] See section 4.1 of chapter 2.

[5] If we restate the zebra case so that this could easily have been the case, then it's no longer clear that

(¬C) If it had not been the case that p, then it would still have been the case that q.

imply corresponding instances of this claim:

(¬E) The fact that q does not establish that p

And, there is good reason to think that this suggestion *is* right. For, the cases in which instances of (¬C) fail to imply instances of (¬E) all seem to be cases in which the antecedent of (¬C) is a *sceptical hypothesis*. Sceptical hypotheses are almost always outlandish claims, which could not easily have been true. Consequently, it's natural to say that (EP_2) is true.

4.5 *From (EP_1) and (EP_2) to (4b')*

Now that we have defended (EP_1) and (EP_2), we can start to argue for (4b'). The first step in our argument is to point out that, if (D) is not true, then the following, counterfactual claim will surely hold:

(i) Even if (D) was true, advocates of EBG would believe EBG, and their reasons for believing it would be exactly the same as they actually are.

This counterfactual claim is equivalent to the following claim (where 'R' refers to the reason for which advocates of EBG believe EBG, whatever that reason is):

(i') If (D) was true, then advocates of EBG would still believe EBG, and would believe it because R.

The next step in our argument is to point out that (i') entails this claim:

(ii) If (D) was true, then it would be the case that R.

---

the zebra case is different from the thermometer case in the way that we are describing.

And, the third step is to point out that, when (ii) is conjoined with (EP_1) and (EP_2), it entails the following claim:

(iii) Under ordinary standards for establishment, the fact that R does not establish that (D) is not true.

The fourth step in our argument points out that, since (as section 4.1 showed) we can infer the falsity of (D) from the truth of EBG, the following claim is true:

(iv) If R doesn't establish that (D) is not true, then R does not establish that EBG is true.

From (iii) and (iv), it follows that, under ordinary standards for establishment, R does not establish that EBG is true. So, when (iii) and (iv) are conjoined with explanatory evidentialism, they entail that, under ordinary standards for knowledge, EBG is not known to be true. Because of this, we can conclude that the following claim is true:

(4b') If (D) is not true, then, under ordinary standards for knowledge, EBG is not known to be true.

4.6 Summary

If the arguments of the last five subsections succeed, then they show that the following claims are both true:

(4a') If (D) is true, then EBG is not true.
(4b') If (D) is not true, then, under ordinary standards for knowledge, EBG is not known to be true.

From these claims, we can infer this claim, which ASO_2 aims to establish:

(4') Under ordinary standards for knowledge, EBG is not known to be true.

If ASO_2 does establish this claim, then all of us have, or can be given, a conclusive epistemic reason to refrain from endorsing EBG. In the next chapter, we will discuss whether it does establish this claim. The upshot of this chapter will be that ASO_2 faces a serious problem. But, in chapter five, we will see that there is a way of resolving this problem.

Chapter Four

A problem for the second anti-sceptical objection

In the last chapter, we introduced a new version of the anti-sceptical objection, which we called *the second anti-sceptical objection*, or ASO_2. The aim of this chapter is to show that there is a serious problem that is faced by advocates of ASO_2. The problem that we will outline is a problem for the second stage of ASO_2, which is an argument for the following claim:

(2') If explanatory evidentialism does not discredit all deity beliefs, then there is no form of evidentialism that discredits all deity beliefs.

In what follows, we will argue that claim (2') is false. We will argue for this by arguing that there is a form of evidentialism that may discredit all deity beliefs even if explanatory evidentialism fails to do so.

The best way of introducing the form of evidentialism which shows claim (2') to be false is to focus on a situation which resembles the job interview case that we described in chapter two. In section 1, we will describe this situation, and will give a knowledge-based explanation of the belief-undermining that takes place in this situation. The explanation that we will give in section 1 is an extended version of the LOK explanation that we defended in chapter two. In sections 2–5, we will defend this extended LOK explanation. Section 6 will then use this explanation to argue that there is a form of evidentialism which may discredit all deity beliefs even if explanatory evidentialism fails to do so.

1. The film director

Two friends, Sarah and Tom, are discussing the politics of film directors. Some way into the discussion, Sarah claims that there aren't any really famous film directors who are left wing. Tom responds to this by claiming that George Lucas is left wing, and the following exchange ensues:

Sarah: I've never heard that George Lucas is left wing. What makes you think that he's left wing?

Tom: Hmm... I can't really remember. I think I read it somewhere.

Sarah: Where did you read it? Was it in one of those socialist newspapers that you're always reading?

Tom: I'm not sure...I suppose it might have been.

Sarah: If it was, then it may well be wrong. The articles in those papers aren't very well researched, and they often misrepresent people's views...

Tom: Yes, I know – but, I'm not sure that I read it in one of those papers. Still, I suppose I may have read it there. I can't be sure that I didn't. And, those papers are pretty unreliable. So, I guess I can't be sure that Lucas really is left wing...

If the situation here is normal, then Sarah's questions and comments will give Tom a conclusive epistemic reason to abandon his active belief that *George Lucas is left wing*. The question that I want to address is: *why* do these questions and comments give Tom a conclusive epistemic reason to abandon this active belief (which we will call his *political belief*)? In the next subsection, I will sketch an answer to this question. Sections 2–5 will then defend this answer, by defending two controversial principles about knowledge.

1.1 The extended LOK explanation

In chapter two, we focused on a case which closely resembles the case of the film director. In this case, a woman, Mary, gives her husband, John, a conclusive epistemic reason to abandon a certain belief by asking him why he has this belief. The events of the case can be summarised in the following way:

(i) First, John tells Mary that he is going to get a certain job.

(ii) Then, Mary asks John why he believes that he is going to get this job.

(iii) John responds by saying that he believes this because his interview for the job went very well.

(iv) Mary then argues that this fact about his interview does not *establish* that he will get the job, and, by doing so, gives John a conclusive epistemic reason to stop actively believing that he will get the job.

One of the aims of chapter two was to explain why Mary's questions and comments undermine John's belief about his job (which we called his *job belief*). In sections 1-5 of this chapter, we argued for the following explanation of this fact (which we called the *lack of knowledge*, or *LOK* explanation):

Mary's questions and comments give John a conclusive epistemic reason to abandon his job belief because they make him aware of the fact that the following claims are both true with respect to him:

(B_j) I believe that *I will get the job* because my interview went well.

($\neg E_j$) The fact that my interview went well does not establish that I will get the job.

John is already aware of the fact that, if (B_j) and ($\neg E_j$) are true with respect to him, then the following claim is also true with respect to him:

($\neg K_j$) I do not know that I will get the job.

So, by making John aware of the fact that (B_j) and ($\neg E_j$) are true with respect to him, Mary makes him aware of the fact that he does not know that he will get the job. And, by making him aware of this fact, she gives him a conclusive epistemic reason to stop actively believing that he will get the job.

It seems clear that, in the case of the film director, Sarah undermines Tom's political belief in roughly the same way as Mary undermines John's job belief. But, it also seems clear that the LOK explanation cannot be applied directly to the case of the film director. For, a direct application would be possible only if Sarah made Tom aware of the fact that corresponding instances of the following claims were true with respect to him:

(B_{lp}) I believe that *George Lucas is left wing* because p.

($\neg E_{lp}$) The fact that p does not establish that George Lucas is left wing.

And, it seems clear that, in the case of the film director, Sarah does *not* make Tom aware of the fact that corresponding instances of these claims are true with respect to him.

Although Sarah does not make Tom aware of the fact that corresponding instances of (B_{lp}) and ($\neg E_{lp}$) *are* true with respect to him, she does make him aware of the fact that the following instances of these claims *might* be true with respect to him:

(B_{ls}) I believe that *George Lucas is left wing* because one of my socialist newspapers said that he was left wing.

($\neg E_{ls}$) The fact that one of my socialist newspapers said that *Lucas is left wing* does not establish that Lucas is left wing.

More exactly: she seems to make him aware of the fact that he does not *know* – and is not in a *position* to know – that (B_{ls}) and ($\neg E_{ls}$) are *not* both true with respect to him. By appealing to this fact about the effect of Sarah's questions and comments, we can construct an explanation of the way in which these questions and comments undermine Tom's political belief. One way of stating this explanation is as follows:

Sarah's questions and comments give Tom a conclusive epistemic reason to abandon his political belief because they make him aware of the fact that *he is not in a position to know that (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him*. Tom is already aware of the fact that, if he is not in a position to know that (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him, then the following claim is true with respect to him:

($\neg K_{ls}$) I do not know that George Lucas is left wing.

So, by making Tom aware of the fact that he is not in a position to know that (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him, Sarah makes him aware of the fact that he does not know that George Lucas is left wing. And, by making him aware of this fact, she gives him a conclusive epistemic reason to stop actively believing that George Lucas is left wing – i.e. to abandon his political belief.

In what follows, we will refer to this as *the extended LOK explanation*. The aim of sections 2–5 is to defend this explanation. In section 2, I will outline a strategy for accomplishing this aim. Sections 3–5 will carry out this strategy by defending two principles about knowledge.

2. Defending the extended LOK explanation

One good way of defending the extended LOK explanation is to defend the following claim:

> (FD) In the case of the film director, Tom is in a position to know that, if he is not in a position to know that (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him, then ($\neg K_{ls}$) is true with respect to him.

If (FD) and the original LOK explanation are both true, then we have good reason to believe the extended LOK explanation. So, since we have already argued that the original LOK explanation is true, we can now defend the extended LOK explanation by defending (FD).

One good way of defending (FD) is to appeal to the following three principles:

> (BEK) We are all aware of the fact that, if an instance of (B) is true, and the corresponding instance of (E) is untrue, then the corresponding instance of (K) is also untrue.[1]

[1] As in chapter two, (B),(E) and (K) refer to the following schematic sentences:

(B) (S believes that p) because q.
(E) The fact that q establishes that p.
(K) S knows that p.

(KKP_K)	We all know that, if a subject, S, knows that p, then S is in a position to know that S knows that p.
(CP_K)	We all know that, if a subject, S, is in a position to know that p, and S is also in a position to know that (if p, then q), then S is in a position to know that q.

The first of these is a principle that we defended in chapter two; the second says that we all know a version of the KK principle, and the third says that we all know that knowledge is closed under known implication. By appealing to these three principles, we can make a strong case for (FD). To see this, note first of all that (BEK) seems to imply the following claim:

(FD_1) In the case of the film director, Tom knows that, if (B_{ls}) and ($\neg E_{ls}$) are both true with respect to him, then ($\neg K_{ls}$) is also true with respect to him.

Next, note that, if (FD_1) is true, then the following claim is very likely also to be true:

(FD_2) In the case of the film director, Tom knows that, if ($\neg K_{ls}$) is not true with respect to him, then (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him.

Thirdly, note that, when (FD_2) is conjoined with (KKP_K), it entails the following claim:

(FD_3) In the case of the film director, Tom is in a position to *know* that he knows that, if ($\neg K_{ls}$) is not true with respect to him, then (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him.

Fourthly: note that, if (FD_3) and (CP_K) are both true, then the following claim is also true:

(FD$_4$) In the case of the film director, Tom is in a position to know that, if he is in a position to know that ($\neg K_{ls}$) is not true with respect to him, then he is also in a position to know that (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him.

Fifthly: note that, if (FD$_4$) and (CP$_K$) are both true, then the following claim is also true:

(FD$_5$) In the case of the film director, Tom is in a position to know that, if he is not in a position to know that (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him, then he is not in a position to know that ($\neg K_{ls}$) is not true with respect to him.

Finally: note that, if (FD$_5$) and (KKP$_K$) are both true, then the following claim is also true:

(FD) In the case of the film director, Tom is in a position to know that, if he is not in a position to know that (B_{ls}) and ($\neg E_{ls}$) are not both true with respect to him, then ($\neg K_{ls}$) is true with respect to him.

The above reasoning shows that, if (BEK), (KKP$_K$) and (CP$_K$) are all true, then (FD) is also likely to be true. We have already defended (BEK); so, by defending (KKP$_K$) and (CP$_K$), we can now defend (FD). In sections 3–5, we will defend (KKP$_K$) and (CP$_K$). If our defence succeeds, then it will give us good reason to believe both (FD) and the extended LOK explanation.

3. Defending (KKP$_K$) and (CP$_K$)

In this section, I will defend (KKP$_K$) and (CP$_K$). I will do so by focusing on a number of puzzles about assertion. Section 3.1 will outline three puzzles of this kind, and section 3.2 will defend a certain solution to these puzzles, which I call *the implication solution*. Section 3.3 will then outline three puzzles which closely resemble the puzzles of section 3.1; and sections 3.4 and 3.5 will argue that the best way of applying the implication solution to these puzzles is to endorse (KKP$_K$) and

(CP_K). If the arguments of sections 3.1–3.5 succeed, then they will give us a strong reason to believe (KKP_K) and (CP_K). In sections 4 and 5, I will argue that this reason is not outweighed by objections to (KKP_K) and (CP_K).

3.1 Three puzzles about assertion

In this section, I will describe three puzzles about assertion. Section 3.2 will then defend a certain solution to these puzzles, which I call *the implication solution*. The puzzles that I will describe are very similar to the puzzles posed by Moorean sentences, like *p but I don't know that p*. The implication solution is inspired by a solution to these Moorean puzzles which is due to Peter Unger, and which was endorsed in chapter two.[2]

3.1.1 The climbing trip

Two friends, A and B are discussing whether to go climbing. B tells A that it would be a bad idea to go climbing, because it is going to snow. A asks B the following question:

(Q_1) How do you know that it is going to snow?

and B gives the following answer:

(A_1) I *don't* know that it is going to snow.

B's answer seems to contradict his earlier assertion that it is going to snow. But, when we look closely at his answer, we can see that it does not in fact contradict this assertion. B's assertion that *it will snow* is perfectly consistent with his assertion that *he doesn't know that it will snow*. So, why does the second assertion seem to contradict the first?

[2] See chapter VI of Unger 1975, esp. pp256-60.

3.1.2 *The cancelled party*

After postponing their climbing trip, A and B discuss their plans for the evening. A suggests going to a party that is being held by their friend C; but, B tells him that C's party has been cancelled. A asks B the following question:

(Q_2) What makes you think that C's party has been cancelled?

and B gives the following answer:

(A_2) I *don't* think that C's party has been cancelled.

Again, B's answer seems to contradict his earlier assertion about the party. And, again, it seems clear, on reflection, that this answer does *not* contradict this assertion. Because of this, we should once again ask why B's answer *seems* to contradict this assertion. Our answer to this question should resemble our answer to the question asked at the end of section 3.1.1.

3.1.3 *The petrol gauge*

Instead of going to C's party, A and B decide to go for a drive. The petrol gauge in A's car is not working, but B assures A that the tank is full. A asks B the following question:

(Q_3) How can you be certain that the tank is full?

and B gives the following answer:

(A_3) I'm *not* certain that the tank is full.

Once again, his answer seems to contradict his earlier assertion; and once again, it seems clear that this apparent contradiction is illusory. To resolve the puzzle posed by this example, we need to explain why B's answer seems to contradict his earlier assertion. Our explanation of this fact should resemble our explanation of the apparent contradictions described in sections 3.1.1 and 3.1.2.

3.2 The implication solution

One natural way of solving the puzzle about the climbing trip is to endorse the following principle about assertion and knowledge:

> (AK) When one asserts that p, one thereby implies that one knows that p.

In the climbing trip example, B asserts, first of all, that *it is going to snow*, and then, that *he does not know that it is going to snow*. If (AK) is true, then, by making the first of these assertions, B implies the negation of the second. So, if (AK) is true, than it is not surprising that B's second assertion seems to contradict his first.

The above solution – which we will call *the implication solution* – is supported by the fact that (AK) is independently plausible. In section 5.2 of chapter two, we defended (AK) by arguing that it can explain a range of puzzling facts about knowledge and assertion. Some of the facts that we mentioned in section 5.2 are quite similar to the fact that we are now using (AK) to explain. But, in spite of this, it seems clear that the discussion of section 5.2 provides us with an independent reason for believing (AK).

Another thing that supports the implication solution is the fact that it generates attractive solutions to both of the other puzzles that we outlined in section 3.1. To see that it generates such solutions, note, first of all, that it is very natural to endorse the following principle about assertion and implication:

> (AI) If, by asserting that p, one implies that q, and it is common knowledge that *if q then r*, then, by asserting that p, one implies that r.

Next, note that the following principles are also very plausible:

(KT_K) It is common knowledge that, if one knows that p, then one thinks that p.

(KC_K) It is common knowledge that, if one knows that p, then one is certain that p.

When (AK) is conjoined with these plausible principles, and with (AI), it entails the following principles:

(AT) When one asserts that p, one implies that one thinks that p.

(AC) When one asserts that p, one implies that one is certain that p.

By appealing to these principles, we can solve the puzzle about the cancelled party, and the puzzle about the petrol gauge in exactly the same way as we just solved the puzzle about the climbing trip.

3.3 Three more puzzles about assertion

In this section, I will describe three more puzzles about assertion.[3] The puzzles that I will describe are very similar to the puzzles that were outlined in section 3.1. The similarity between the two sets of puzzles strongly suggests that the implication solution should be applied to both. In sections 3.4 and 3.5, I will argue that the best way of applying the implication solution to the puzzles that are outlined in this section is to endorse (KKP_K) and (CP_K).

3.3.1 Another climbing trip

Two friends, A and B are discussing whether to go climbing. B tells A that it would be a bad idea to go climbing, because it is going to snow. A asks B the following question:

[3] The puzzles that I will describe correspond to the puzzles posed by iterated Moorean sentences, like 'p, but I don't know whether I know that p', 'p, but I don't know whether I think that p' and so on. For discussion of such sentences, see Hambourger 1987:252 and Sorensen 2000.

(Q_1) How do you know that it is going to snow?

and B gives the following answer:

(A_{1K}) I don't know whether I *do* know that it is going to snow.

B's answer seems to contradict his earlier assertion that it is going to snow. But, when we look closely at his answer, we can see that it does not in fact contradict this assertion. B's assertion that *it will snow* is perfectly consistent with his assertion that *he doesn't know whether he knows that it will snow*. So, why does the second assertion seem to contradict the first?

3.3.2 *Another cancelled party*

After postponing their climbing trip, A and B discuss their plans for the evening. A suggests going to a party that is being held by their friend C; but, B tells him that C's party has been cancelled. A asks B the following question:

(Q_2) What makes you think that C's party has been cancelled?

and B gives the following answer:

(A_{2K}) I don't know whether I *do* think that C's party has been cancelled.

Again, B's answer seems to contradict his earlier assertion about the party. And, again, it seems clear, on reflection, that this answer does *not* contradict this assertion. Because of this, we should again ask why B's answer *seems* to contradict this assertion. Our answer to this question should resemble our answer to the question asked at the end of section 3.3.2.

3.3.3 *Another petrol gauge*

Instead of going to C's party, A and B decide to go for a drive. The petrol gauge in A's car is not working, but B assures A that the tank is full. A asks B the following question:

(Q_3) How can you be certain that the tank is full?

and B gives the following answer:

(A_{3K}) I don't know whether I *am* certain that the tank is full.

Again, his answer seems to contradict his earlier assertion; and again, it seems clear that this apparent contradiction is illusory. To resolve the puzzle posed by this example, we need to explain why B's answer seems to contradict his earlier assertion. Our explanation of this fact should resemble our explanation of the apparent contradictions described in sections 3.3.1 and 3.3.2.

3.4 *Extending the implication solution*

The three puzzles just outlined are very similar to the three puzzles outlined in section 3.1. So, there is reason to think that the implication solution can be applied to the puzzles that we have just outlined. In this section, I will argue that the best way of applying this solution to the new climbing trip puzzle is to endorse (KKP_K). The next section will argue that the best way of applying the solution to the other puzzles that we have just outlined is to endorse (KKP_K) and (CP_K).

The most straightforward way of applying the implication solution to the new climbing trip puzzle is to claim that the following principle is true:

(KK_K) It is common knowledge that, if one knows that p, then one also knows that one knows that p.

When this principle is conjoined with the following pair of principles, which are at the heart of the implication solution, it generates a solution to the new climbing trip puzzle:

(AK) When one asserts that p, one thereby implies that one knows that p.

(AI) If, by asserting that p, one implies that q, and it is common knowledge that *if q then r*, then, by asserting that p, one implies that r.

For, when it is conjoined with (AK) and (AI), (KK_K) entails that the following principle is true:

(AKK) When one asserts that p, one thereby implies that one knows that one knows that p.

And, if this principle is true, then it is not surprising that, when B asserts that *he does not know whether he knows that it is going to snow*, his assertion seems to contradict his earlier assertion that *it is going to snow*.

The main problem with this solution to the new climbing trip puzzle is that there are strong reasons for rejecting (KK_K). One such reason is that, if (KK_K) is true, then possession of knowledge is restricted to those who have the concept of knowledge. Another is that, if (KK_K) is true, then one cannot know a proposition without knowing an infinite number of other propositions, many of which are propositions that none of us have ever entertained.[4] Objections of this kind suggest that, if we want to apply the implication solution to the new climbing trip puzzle, then we should appeal to a weaker principle that (KK_K). One natural way of weakening (KK_K) is to rewrite it as follows:

(KKP_K) It is common knowledge that, if one knows that p, then one is *in a position* to know that one knows that p.

If the notion of being *in a position* to know is understood in an appropriate way, then this new principle will be invulnerable to the objections that are usually levelled against principles like (KK_K). One appropriate way of understanding it is as a *counterfactual* notion, on which S is in a position to know that p iff something like the following counterfactual is true:

[4] For a statement of these and other problems for (KK_K) see Sorensen 1988:242.

(CF) If S had the concept of knowledge, and S was to reflect about whether p, then S would come to know that p.

If we understand it in this way, then we can reasonably endorse the following principle:

(PK_K) It is common knowledge that, if S is in a position to know that p, then S does not know that *S does not know whether p.*

And, if we conjoin this principle with (AK), (AI) and (KKP_K), then we can generate another solution to the new climbing trip puzzle.

To see this, note first of all that, when (KKP_K) is conjoined with (AK) and (AI), it entails the following principle about assertion:

(AKKP) When a subject, S, asserts that p, S implies that he is in a position to know that he knows that p.

Next: note that, when this principle is conjoined with (AI) and (PK_K), it entails another principle about assertion:

(AKKK) When a subject, S, asserts that p, S implies that he does not know that (he does not know whether he knows that p).

Finally: note (i) that, if (AKKK) is true, then when B asserts that *it is going to snow*, he implies that he does not know that (he does not know whether he knows that it is going to snow), and (ii) that, if (AK) is true, then when B asserts that *he does not know whether he knows that it is going to snow*, he implies that he *does* know that (he does not know whether he knows that it is going to snow). By noting this, we can see that, if (AKKK) and (AK) are both true, then they will explain why B's second assertion seems to contradict his first. And once we have seen this, it is natural to think that the best way of applying the implication solution to the new climbing trip puzzle is to endorse (KKP_K).

3.5 *The other puzzles*

If the best way of applying the implication solution to the new climbing trip puzzle is to endorse (KKP_K), then we have reason to believe (KKP_K). In this section, I will argue that we have similar reason to believe the conjunction of (KKP_K) and (CP_K). I will argue for this by focusing on the other two puzzles that we outlined in section 3.3. In what follows, I will argue that (KKP_K) and (CP_K) jointly provide us with an attractive way of applying the implication solution to these puzzles.

In the last section, we saw that, when (KKP_K) is conjoined with (AK) and (AI), it entails the following principle about assertion:

(AKKP) When a subject, S, asserts that p, S implies that he is in a position to know that he knows that p.

In what follows, I will argue that, by conjoining this principle about assertion with (CP_K), and with some other plausible principles, we can generate an attractive solution to the puzzles that we outlined in sections 3.3.2 and 3.3.3.

The first step in my argument is to draw out some consequences of (CP_K). According to (CP_K), it is common knowledge that, if one is in a position to know that p, and one is in a position to know that *if p then q*, then one is in a position to know that q. If this really is common knowledge, then we have reason to accept the following principles:

(KKT_K) It is common knowledge that, if one is in a position to know that one knows that p, then one is in a position to know that one *thinks* that p.

(KKC_K) It is common knowledge that, if one is in a position to know that one knows that p, then one is in a position to know that one is *certain* that p.

For, as we saw earlier, the following principles are both very plausible:

(KT_K)	It is common knowledge that, if one knows that p, then one thinks that p.
(KC_K)	It is common knowledge that, if one knows that p, then one is certain that p.

And, if these principles and (CP_K) are all true, then it is natural to think that (KKT_K) and (KKC_K) are also true.

When (KKT_K) and (KKC_K) are conjoined with (AKKP) and (AI), they entail the following principles about assertion:

(AKTP)	When a subject, S, asserts that p, S implies that he is in a position to know that he thinks that p.
(AKCP)	When a subject, S, asserts that p, S implies that he is in a position to know that he is certain that p.

And, when these principles are conjoined with (AI) and with the following principle (which we defended earlier):

(PK_K)	It is common knowledge that, if S is in a position to know that p, then S does not know that he does *not* know whether p.

they entail these two principles, which can be used to solve the puzzles that were outlined in sections 3.3.2 and 3.3.3:

(AKTK)	When a subject, S, asserts that p, S implies that he does not know that (he does not know whether he thinks that p).
(AKCK)	When a subject, S, asserts that p, S implies that he does not know that (he does not know whether he is certain that p).

To see that these principles can be used for this purpose, we need only recap on the details of the two puzzles. In the puzzle of section 3.3.2, B asserts first of all, that *C's party has been cancelled*, and then, that *he does not know whether he thinks that C's* *party has been cancelled.* If (AKTK) is true, then, by making his first assertion, B implies that he does not know that (he does not know whether he thinks that C's party has been cancelled); and, if (AK) is true, then, by making his second assertion, B implies that he *does* know that (he does not know whether he thinks that C's party has been cancelled); so, if (AKTK) and (AK) are both true, then it is not surprising that B's second assertion seems to contradict his first.

In the puzzle of section 3.3.3, B asserts first of all, that *the tank is full*, and then, that *he does not know whether he is certain that the tank is full.* If (AKCK) is true, then, by making his first assertion, B implies that he does not know that (he does not know whether he is certain that the tank is full); and, if (AK) is true, then, by making his second assertion, B implies that he *does* know that (he does not know whether he is certain that the tank is full); so, if (AKCK) and (AK) are both true, then it is again unsurprising that B's second assertion seems to contradict his first.

The upshot of the last two sections is that (KKP_K) and (CP_K) jointly provide us with an attractive way of applying the implication solution to the puzzles that were outlined in section 3.3. If the arguments of these sections succeed, then they give us a reason to believe (KKP_K) and (CP_K). Some may claim that this reason is outweighed by certain objections to (KKP_K) and (CP_K). In the next two sections, I will attack this claim, by attacking what I take to be the best objections to (KKP_K) and (CP_K).

4. An objection to (CP_K)

In this section, I will outline and evaluate what I take to be the best objection to the following principle:

> (CP_K) We all know that, if a subject, S, is in a position to know that p, and S is also in a position to know that (if p, then q), then S is in a position to know that q.

The objection that I will discuss is closely based on an objection to the closure principle that Fred Dretske gives in his paper 'Epistemic Operators'.[5] One good way of stating this objection is to appeal to an example that Dretske describes in this paper.[6] In section 4.1, I will use this example to state the objection, and in section 4.2, I will argue that the objection fails.

4.1 A Dretskean objection

During a visit to my local zoo, I stop outside the zebra enclosure. Inside the enclosure are several zebras, and a sign which says that these animals are zebras. Do I know that the animals in the enclosure are zebras? If circumstances are normal, then it seems clear that I do.

Now consider the (admittedly bizarre) claim that the animals in the zebra enclosure are mules that have been cleverly disguised to look like zebras (henceforth: cleverly disguised mules). It seems clear that this claim is very implausible, and that I know it to be very implausible. But, in spite of this, it does not seem right to say that I know this claim to be *false*. To *know* that this claim was false, I would have to investigate whether the animals in the zebra enclosure have been disguised; and, if circumstances are normal, then I will not have conducted any investigation of this kind. Consequently, it is natural to say that, if circumstances are normal, then I do not know – and am not in a position to know – that the animals in the enclosure are not cleverly disguised mules. But, if this is so, then (CP_K) is surely false – for, as we have already seen, I *do* seem to know that the animals in the enclosure are zebras, and also seem to be in a position to know that, if they are zebras, then they are not cleverly disguised mules.

4.2 A response to this objection

The objection of section 4.1 trades on the fact that, in the Dretskean example that we have just described, each of the following claims is plausible:

(D_1) I know that the animals in the zebra enclosure are zebras.

---

[5] Dretske 1970.
[6] Dretske 1970:1015-6.

(D_2) I am in a position to know that, if the animals in the zebra enclosure are zebras, then they are not cleverly disguised mules.

(D_3) I do not know, and am not in a position to know that the animals in the zebra enclosure are not cleverly disguised mules.

If these three claims are all true, then (CP_K) is false. The objection of section 4.1 makes each claim seem true, and so makes it seem likely that (CP_K) is false. However: the objection loses some of its force when we focus on the following claim, which is generated by conjoining (D_1) and (D_3):

($D_{1\&3}$) I know that the animals in the enclosure are zebras, but I do not know, and am not in a position to know that they are not cleverly disguised mules.

Although advocates of section 4.1's objection make (D_1) and (D_3) both seem plausible, they do not make ($D_{1\&3}$) seem plausible. When we reflect on whether ($D_{1\&3}$) is true in the example of section 4.1, our inclination is to say that it is not. Advocates of section 4.1's objection are committed to saying that ($D_{1\&3}$) *is* true in this example; so, the implausibility of this claim presents such advocates with a problem. In what follows, I will develop this problem into a rebuttal of section 4.1's objection. I will do so by arguing that the best way of explaining the difference in plausibility between the individual claims (D_1) and (D_3) and the conjunctive claim ($D_{1\&3}$) is to endorse a form of epistemic contextualism.

According to epistemic contextualists, the truth-conditions of our knowledge attributions vary with the context in which those attributions are made. By endorsing a certain form of epistemic contextualism, we can explain why (D_1) and (D_3) are plausible when considered individually, but implausible when considered in conjunction. The contextualist theory that I have in mind is a theory that is due to Keith DeRose.[7] The core claim of this theory is that the truth-conditions of our knowledge attributions are governed by the following *rule of sensitivity* (where S's belief that p is said to be *sensitive* iff S would not have believed that p, if it had not

[7] This theory is defended in DeRose 1995.

been the case that p):

(ROS) When it is asserted that some subject, S, knows (or doesn't know) some proposition, P, the standards for knowledge tend to be raised in such a way as to require S's belief in that particular P to be sensitive for it to count as knowledge.[8]

By appealing to (ROS), we can construct an explanation of why (D_1) and (D_3) seem true individually, but false in conjunction. According to this explanation – which we will call *the contextualist explanation* – (D_1) and (D_3) seem true individually because each is evaluated under different standards for knowledge. More exactly: the contextualist explanation says (a) that, when (D_1) is asserted, at the start of section 4.1, the standards for knowledge are such as to make (D_1) true, and (D_3) false, and (b) that, when (D_3) is asserted, at the end of section 4.1, the standards for knowledge are raised in such a way as to make (D_1) false, and (D_3) true.

To see that (ROS) supports the contextualist explanation, we need to focus on the following pair of counterfactuals:

(C_z) If the animals in the enclosure had not been zebras, then I would not have believed that they were zebras.

(C_m) If the animals in the enclosure had been cleverly disguised mules, then I would not have believed that they were not cleverly disguised mules.

It seems clear that, in the situation described in section 4.1, (C_z) is true, and (C_m) false. Consequently, it seems clear that, in this situation, I sensitively believe that the animals in the enclosure are zebras, but do not sensitively believe that they are not cleverly disguised mules. Because of this, (ROS) entails that, when (D_3) is asserted, at the end of section 4.1, the standards for knowledge are raised in such a way as to make (D_3) true. If the standards for knowledge are raised in this way, then it is natural to suppose that the contextualist explanation is correct.

[8] Cf. DeRose 1995:36.

If the contextualist explanation is correct, then the objection of section 4.1 fails. But, it is not yet clear that the contextualist explanation is correct. In the remainder of this section, I will give three reasons for believing this explanation. Once these reasons have been outlined, it should be clear that the objection of section 4.1 fails.

If the contextualist explanation is correct, then (D_1) should cease to seem plausible, after (D_3) has been asserted. On reflection, it seems clear that (D_1) *does* cease to seem plausible, after this assertion has been made. Once it has been asserted that *I do not know that the animals in the zebra enclosure are not cleverly painted mules*, it no longer seems right to say that I *do* know that these animals are not zebras. The fact that this no longer seems right is a reason to endorse the contextualist explanation.

Another reason for endorsing the contextualist explanation is that (ROS), which supports this explanation, is independently plausible. In his 1995, DeRose shows that, by appealing to (ROS), we can generate attractive solutions to a range of sceptical paradoxes; and, in his 1996, he shows that this rule can also help us to understand certain puzzling intuitions about lotteries. The explanatory utility of (ROS) gives us reason to believe that (ROS) is true. By doing this, it gives us reason to endorse the contextualist explanation, which is supported by (ROS).

A third reason for endorsing the contextualist explanation is generated by the arguments for epistemic contextualism that we gave in section 4.3 of chapter two. If these arguments succeed, then they provide further support for (ROS). For, if they succeed, then they show that the following contextualist claim is true:

> (CCK) When it is asserted (or implied) that an instance of (B) is true and that the corresponding instance of (K) is true, the standards for knowledge tend to be raised in such a way as to ensure that, if those instances of (B) and (K) are true then the corresponding instance of (C) is true.[9]

[9] Here, (B), (K) and (C) refer to the following schematic sentences:

(B) (S believes that p) because q.

And, if this claim is true, then it is very likely that (ROS) is also true – for, corresponding instances of (B) and (K) tend to be true when, and only when, the relevant instance of (B) refers to a belief that is sensitive.

5. An objection to (KKP_K)

The last section rebutted the best objection to (CP_K). This section will rebut the best objection to the following claim:

(KKP_K) We all know that, if a subject, S, knows that p, then S is in a position to know that S knows that p.

The objection that I will discuss is closely based on an objection to the KK principle that Timothy Williamson gives in his paper 'Inexact Knowledge'.[10] One good way of stating this objection is to appeal to an example that Williamson describes in this paper.[11] In sections 5.1 and 5.2, I will describe this example, and will explain how it threatens (KKP_K). Sections 5.3 and 5.4 will then argue that this objection fails.

5.1 The distant tree

As I look out of my window, I can see a distant tree. Wondering how tall it is, I try to judge the height of the tree. I am not able to judge its height to the nearest inch; so, although the tree is 600 inches tall, I do not come to *know* that it is 600 inches tall. However: I do seem to learn some things about the height of the tree. One thing that I seem to learn is that the tree is not 60 inches tall. Another is that it is not 6000 inches tall. And, a third is that its height is somewhere between 60 and 6000 inches.

As I am in a philosophical mood, I start to reflect about my knowledge of the tree. My reflections can be set down in the following way:

(K) S knows that p.
(C) If it was not the case that p, then it would not be the case that q.

[10] Williamson 1992.
[11] Williamson 1992:217-9.

My eyesight and my ability to judge heights are limited. Consequently, there are certain things that I cannot come to know about the height of the tree by just looking at it from my window. If the tree is m inches tall, then I cannot come to know, by just looking at it from my window, that it is m inches tall. And, if the tree is m+1 inches tall, or m-1 inches tall, then I cannot come to know, by just looking at it from my window, that it is *not* m inches tall. Everything that I now know about the height of the tree is something that I have come to know by just looking at the tree from my window. So, if the tree is m inches tall, then I do not know that it is m inches tall; and, if it is m+1 inches tall, or m-1 inches tall, then I do not know that it is *not* m inches tall.

As a result of these reflections, I come to believe that the following principle is true:

(M) For any natural number m: if the tree is m+1 inches tall, then I do not know that it is not m inches tall.

I then start to make some inferences. The first thing that I do is to deduce, from (M), that the following principle is true:

(M') For any natural number m: if I know that the tree is not m inches tall, then it is not m+1 inches tall.

I then deduce, from (M'), that the following conditional is true:

(K_{60}) If I know that the tree is not 60 inches tall, then the tree is not 61 inches tall.

After this, I deduce, from (M'), that this conditional is true:

(K_{61}) If I know that the tree is not 61 inches tall, then the tree is not 62 inches tall.

and that this conditional is true:

(K_{62}) If I know that the tree is not 62 inches tall, then the tree is not 63 inches tall.

and so on, until I have deduced that the following conditional is true:

(K_{599}) If I know that the tree is not 599 inches tall, then the tree is not 600 inches tall.

I then reflect upon whether the antecedent of (K_{60}) is true. After concluding that it *is* true, I deduce, from this conclusion, and from (K_{60}), that the consequent of (K_{60}) is also true. Having deduced this, I then ask myself whether the antecedent of (K_{61}) is true. After concluding that it *is* true, I deduce, from this conclusion, and from (K_{61}), that the consequent of (K_{61}) is also true, and then go on to consider the antecedent of (K_{62}). After repeating this process many times, I finally come to reflect upon whether the antecedent of (K_{599}) is true. I conclude that it *is* true, and then deduce that the consequent of (K_{599}) is also true – i.e. that the tree is not 600 inches tall. It seems clear that, by deducing this, I cannot come to *know* that the tree is not 600 inches tall – for, as we stipulated at the start of this subsection, the tree *is* 600 inches tall. However, it can be argued that, if (KKP_K) is true, then, by deducing that the consequent of (K_{599}) is true, I *do* come to know that the tree is not 600 inches tall. In section 5.2, I will outline an argument for this claim.[12] If this argument succeeds, then it shows that (KKP_K) is not true. In sections 5.3 and 5.4, I will show that the argument does not succeed.

5.2 The argument

In the example of the distant tree, some reflections on the limitations of my eyesight, and my height-judging abilities, lead me to believe that the following principle is true:

(M) For any natural number m: if the tree is m+1 inches tall, then I do not know that it is not m inches tall.

[12] This argument is derived from the argument in sec 5.1 of Williamson 2000.

It seems clear that, as a result of these reflections, I come to *know* that principle (M) is true. It also seems clear that, if I come to know this principle, then I come to know every proposition that I properly deduce from this principle. In the example of the distant tree, I properly deduce, from (M), every instance of the following conditional in which 'm' is replaced by a numeral between '60' and '599' :

(K_m) If I know that the tree is not m inches tall, then the tree is not m+1 inches tall.

Consequently, we can conclude that, in the example of the distant tree, I come to know every instance of (K_m) in which 'm' is replaced by a numeral in this range.

One instance of (K_m) that I come to know, in the example of the distant tree, is the following:

(K_{60}) If I know that the tree is not 60 inches tall, then the tree is not 61 inches tall.

After I have come to know (K_{60}), I reflect upon whether its antecedent is true, and conclude that it *is* true. Do I thereby come to *know* that the antecedent of (K_{60}) is true? If (KKP_K) is true, then I do. For, it seems clear that, in the example of the distant tree, the antecedent of (K_{60}) is true. And, (KKP_K) entails that, if the antecedent of (K_{60}) is true, and I reflect upon whether it is true, then I will thereby come to *know* that it is true.

After I have concluded that the antecedent of (K_{60}) is true, I deduce from this, and from (K_{60}), that the consequent of (K_{60}) is true. If (KKP_K) is true, then, by doing this, I properly deduce, from propositions that I know, that the consequent of (K_{60}) is true. Proper deduction of this kind seems to preserve knowledge; so, if (KKP_K) is true, then, by making this deduction, I come to know that the consequent of (K_{60}) is true.

Once I have deduced that the consequent of (K_{60}) is true, I reflect upon whether the antecedent of the following claim is true, and conclude that it *is* true:

(K_{61}) If I know that the tree is not 61 inches tall, then the tree is not 62 inches tall.

Do I thereby come to know that the antecedent of (K_{61}) is true? If (KKP_K) is true then I do. For, if (KKP_K) is true, then, by deducing that the consequent of (K_{60}) is true, I have come to *know* that the consequent of (K_{60}) is true, and have thus made the antecedent of (K_{61}) true. (KKP_K) entails that, if the antecedent of (K_{61}) is true, and I reflect upon whether it is true, then I will thereby come to know that it is true. So, if (KKP_K) is true, then, by reflecting on whether the antecedent of (K_{61}) is true, I come to know that the antecedent of (K_{61}) is true.

After I have concluded that the antecedent of (K_{61}) is true, I deduce, from this and from (K_{61}), that the consequent of (K_{61}) is true. I then reflect upon whether the antecedent of (K_{62}) is true, and, after concluding that it is true, deduce that the consequent of (K_{62}) is also true. After repeating this process many times, I am finally led to deduce that the consequent of (K_{599}) is true – i.e. that the tree is not 600 inches tall. If (KKP_K) is true, then, by deducing this, I come to *know* that the tree is not 600 inches tall. But, it seems clear that, by deducing this, I do *not* come to know that the tree is not 600 inches tall; and, consequently, we can conclude that (KKP_K) is not true.

5.3 *Evaluating the argument*

The argument of section 5.2 seems to have four premises. The first premise is that, in the case of the distant tree, I come to know that the following principle is true:

(M) For any natural number m: if the tree is m+1 inches tall, then I do not know that it is not m inches tall.

The second premise is that, in the case of the distant tree, I come to know everything that I deduce, properly, from propositions that I know. The third premise is that, in the case of the distant tree, I know that the tree is not 60 inches tall. And, the fourth premise is that, in the case of the distant tree, I do *not* come to know that the tree is not 600 inches tall.

It seems clear that, if these premises are all true, then the argument of section 5.2 succeeds. Consequently, we can evaluate the argument by evaluating each of these premises. The fourth premise of the argument is very hard to deny; for, it seems clear (a) that, in the case of the distant tree, the tree is 600 inches tall, and (b) that the case of the distant tree is a *possible* case – and hence, not a case in which the tree both is, and is not 600 inches tall. The third premise is also hard to deny; for, it seems clear that, if the tree is 600 inches tall, then I will learn, by looking at the tree, that it is not 60 inches tall. Some may try to attack the second premise by appealing to Dretskean objections to deductive closure; but, as we have already seen, there is good reason to think that such objections fail. So, the only premise of the argument that we seem to have a chance of rebutting is the first premise, which says that, in the case of the distant tree, I come to know that the following principle is true:

(M) For any natural number m: if the tree is m+1 inches tall, then I do not know that it is not m inches tall.

One good way of attacking this premise is to ask why it is that I come to believe (M), in the case of the distant tree. It is natural to say that, in the case of the distant tree, I come to believe (M) because I recognise that (M) is entailed by the following pair of principles:

(M_1) For any natural number m: if the tree is m+1 inches tall, then I cannot come to know, by just looking at the tree, that the tree is not m inches tall.

(M_2) Everything that I know about the height of the tree is something that I have come to know by just looking at the tree.

But, it is not clear why I believe these principles, in the case of the distant tree. In particular: it is not clear why I believe (M_1), in this case. In the remainder of this subsection, I will argue that, in the case of the distant tree, I believe (M_1) because I believe the following principle:

(M_1*) For any natural number m: if the tree is m+1 inches tall, and I come to believe, by just looking at the tree, that *the tree is not m inches tall*, then my belief that *the tree is not m inches tall* is based on insensitive evidence.

Section 5.4 will then argue that, if my belief in (M_1) is explained by a belief in (M_1*), then the argument of section 5.2 fails.

Before we can argue that my belief in (M_1) is explained by a belief in (M_1*), we need to explain what it is for a belief to be *based on insensitive evidence*. One good way of clarifying the meaning of this phrase is to focus on the following schematic sentences:

(B_t) I believe that *the tree is not m inches tall* because q.

(C_t) If it had not been the case that *the tree is not m inches tall*, then it would not have been the case that q.

When I say that *my belief that the tree is not m inches tall is based on insensitive evidence*, what I mean is that *an instance of (B_t) is true, and the corresponding instance of (C_t) is false*. Once this is clarified, it becomes fairly clear that, in the example of the distant tree, (M_1*) is true. To see this, note first of all that, if I come to believe, by just looking at the tree, that *the tree is not m inches tall*, then something like the following instance of (B_t) will be true:

(B_{tl}) I believe that *the tree is not m inches tall* because it looks to me as if the tree is not m inches tall.

Next, note that, if, in the example of the distant tree, it looks to me as if *the tree is not m inches tall* when the tree is in fact m+1 inches tall, then the following claim will surely be true:

($\neg C_{tl}$) If the tree had been m inches tall, then it might still have looked to me as if the tree was not m inches tall.

If $(\neg C_{tl})$ is true, then the instance of (C_t) which corresponds to (B_{tl}) will be false. Because of this, it is plausible to suppose that, in the example of the distant tree, $(M_1{*})$ is true.

Although it is plausible to suppose that, in the example of the distant tree, $(M_1{*})$ is true, it is not yet clear that, in the example of the distant tree, my belief in (M_1) is explained by my belief in $(M_1{*})$. In what follows, I will try to make this clear. I will do so in two stages. First of all, I will use a psychological generalisation that I defended in chapter two to argue that, if $(M_1{*})$ seems true, then (M_1) will also seem true. Then, I will argue that alterations to the distant tree case which make (M_1) seem false also make $(M_1{*})$ seem false.

In chapter two, I defended a psychological generalisation about knowledge. If we appeal to the following schematic sentences:

(B) (S believes that p) because q.

(C) If it was not the case that p, then it would not be the case that q.

(K) S knows that p.

then we can state this generalisation in the following way:

(BCK_j) When we judge that an instance of (B) is true, and that the corresponding instance of (C) is untrue, we tend also to judge that the corresponding instance of (K) is untrue.

If this generalisation is true, and $(M_1{*})$ seems true, then (M_1) will also seem true. To see this, note first of all that $(M_1{*})$ can be restated in the following way (where '(B_t)' and '(C_t)' are defined as before):

$(M_1{*}')$ For any natural number m: if the tree is m+1 inches tall, and I come to believe, by just looking at the tree, that *the tree is not m inches tall*, then an instance of (B_t) is true, and the corresponding instance of (C_t) is untrue.

Next, note that, if (BCK_j) is true, and the consequent of (M_1*') seems true, then the following claim will seem untrue:

(K_t) I know that *the tree is not m inches tall.*

From this, it follows that, if (BCK_j) is true, and (M_1*) seems true, then the following claim will also seem true:

(M_1') For any natural number m: if the tree is m+1 inches tall, and I come to believe, by just looking at the tree, that *the tree is not m inches tall*, then I do not know that the tree is not m inches tall.

If (M_1') seems true, then (M_1) will also seem true; so, by appealing to chapter one's arguments for (BCK_j), we can argue that, if (M_1*) seems true, then (M_1) will also seem true.

The above considerations show that, in the case of the distant tree, my belief in (M_1) *could* be explained by a belief in (M_1*). But, they do not show that, in this case, my belief in (M_1) is *in fact* explained in this way. To argue for this stronger conclusion, we need to modify the case of the distant tree. One good way of modifying it is to imagine that, in this case, I *am* able visually to judge the height of distant trees to the nearest inch. When the case is modified in this way, we cease to be inclined to say that, in this case, (M_1) is true. Interestingly, though, we also cease to be inclined to say that, in this case, (M_1*) is true. To see this, recall that the plausibility of (M_1*) is due to the fact that, in the example of the distant tree, the antecedent of (M_1*)'s conditional seems to imply both of the following claims:

(B_{tl}) I believe that *the tree is not m inches tall* because it looks to me as if the tree is not m inches tall.

($\neg C_{tl}$) If it had not been the case that *the tree is not m inches tall* (i.e. if the tree had been m inches tall), then it might still have been the case that *it looks to me as if the tree is not m inches tall* (i.e. it might still have looked to me as if the tree was not m inches tall).

When the case of the distant tree is altered in the way that we just described, it no longer seems right to say that the antecedent of (M_1*)'s conditional implies both of these claims. For, when the case is altered in this way, it no longer seems right to say that, if it looks to me as if *the tree is not m inches tall* when the tree is in fact m+1 inches tall, then ($\neg C_{tl}$) is true.

The above considerations give us reason to believe that, in the example of the distant tree, my belief in (M_1) is *in fact* explained by my belief in (M_1*). In the next section, I will argue that, if my belief in (M_1) is explained in this way, then the argument of section 5.2 fails.

5.4 Contextualism and knowledge of (M_1)

Suppose that, in the example of the distant tree, my belief in (M_1) is explained by my belief in (M_1*). From this supposition, it clearly follows that, in the example of the distant tree, my belief in (M_1) constitutes knowledge only if (M_1) is implied by (M_1*). Is (M_1) implied by (M_1*)? That depends on whether the following claim is true (where '(B_t)', '(C_t)' and '(K_t)' are used in the same way as before) :

> (BCK_t) If an instance of (B_t) is true, and the corresponding instance of (C_t) is not true, then the corresponding instance of (K_t) is not true.

In section 4.3 of chapter two, we discussed whether claims like (BCK_t) are true. The upshot of that section was that, although such claims are not invariably true, the standards for knowledge can be raised in such a way as to make them true. Consequently, we should now conclude that, although (M_1*) does not invariably imply (M_1), the standards for knowledge can sometimes be raised in such a way as to make (M_1*) imply (M_1). And, this should lead us to conclude that, although the following claim is not invariably true, the standards for knowledge can sometimes be raised in such a way as to make it true:

> (MK) In the example of the distant tree, my belief in (M_1) constitutes knowledge.

The first premise of the argument of section 5.2 says that, in the example of the distant tree, my belief in principle (M) constitutes knowledge. In the example of the distant tree, I believe (M) because I believe (M_1); so, it seems clear that, if (MK) is not true, then the first premise of the argument of section 5.2 also is not true. We have just argued that (MK) is not invariably true; so, we can now conclude that the first premise of the argument of section 5.2 also is not invariably true. However: we have also argued that the standards for knowledge can sometimes be raised in such a way as to make (MK) true. So, we should conclude that the same may be true with respect to the first premise of the argument in section 5.2.

The above remarks show that the first premise of the argument of section 5.2 is true only in contexts in which the standards for knowledge have been raised. By doing so, they provide advocates of (KKP_K) with a way of responding to the argument of section 5.2. To state this response, we need to focus on the third premise of this argument, which says that, *in the example of the distant tree, I know that the tree is not 60 inches tall.* The gist of the response is that there may be *no* context in which the first premise and the third premise of section 5.2's argument are both true.

Advocates of this response do not need to *show* that there is no context in which the first and third premises of section 5.2's argument are both true. All that they need to do is to point out that there *may* be no context of this kind. If there is no context of this kind, then the argument of section 5.2 fails. So, by pointing out that there may be no such context, they present advocates of this argument with a serious challenge.

In order to respond to this challenge, advocates of section 5.2's argument must show that there *is* some context in which the first and third premises of this argument are both true. It is not easy to see how they could accomplish this goal. Of course, it may turn out that the goal can be accomplished. But, until there is some positive reason for thinking that this is so, it seems fair to conclude that the argument of section 5.2 fails.

5.5 Summary

The aim of this section, and of the previous two, has been to defend the following pair of claims:

(KKP_K)	We all know that, if a subject, S, knows that p, then S is in a position to know that S knows that p.
(CP_K)	We all know that, if a subject, S, is in a position to know that p, and S is also in a position to know that (if p, then q), then S is in a position to know that q.

In section 3, we gave an argument for this pair of claims, and in this section, and the last section, we showed that the force of this argument is not diminished by the best objections to these claims.

In section 2, we argued that, by defending (KKP_K) and (CP_K), we can also defend a certain explanation, which we called *the extended LOK explanation*. If the arguments of sections 2–5 all succeed, then they give us good reason to endorse the extended LOK explanation. In section 6, we will use the extended LOK explanation to generate a problem for the second anti-sceptical objection (which we are calling ASO_2). The gist of this problem is that the second stage of ASO_2 fails.

6. The extended LOK explanation and ASO_2

Before we can show that the extended LOK explanation generates a problem for ASO_2, we must trace out some of the consequences of this explanation. In sections 6.1 and 6.2, we will argue that, if the extended LOK explanation is correct, then we are all in a position to know the following doctrine:

> (E_{EP}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S is not in a position to know that his belief in P is explained by establishing evidence, then S's belief in P does not constitute knowledge.

Section 6.3 will argue that this doctrine is a form of evidentialism, and section 6.4 will argue that, by appealing to this form of evidentialism, we can generate a problem for ASO_2.

6.1 *Generalising the extended LOK explanation*

If the extended LOK explanation is correct, then another, more general explanation is also correct. To see this, recall the case of the film director, on which the extended LOK explanation is focused. The case of the film director conforms to the following pattern (which we will call *the beta pattern*):

(i) First, person A tells person B that p.

(ii) Then, B asks A *why he believes* that p.

(iii) A responds by saying, sincerely, that he does not know why he believes that p.

(iv) B then points out (a) that, for all A knows, he believes that p because q, and (b) that the fact that q does not *show*, or *establish* that p.

In virtually every case that conforms to this pattern, B's questions and comments seem to give A a conclusive epistemic reason to stop actively believing that p.

Why is it that, in virtually every beta-pattern case, B's questions and comments undermine A's active belief that p? If the extended LOK explanation is correct, then we should surely explain this fact in the following way:

In virtually every beta-pattern case, B's questions and comments make A aware of the fact that he is not in a position to know that the following claims are not both true with respect to him:

(B_p) (I believe that p) because q.

($\neg E_p$) The fact that q does not establish that p.

And, in virtually every case of this kind, A is already aware of the fact that, if he is not in a position to know that (B_p) and ($\neg E_p$) are not both true with respect to him, then the following claim is true with respect to him:

($\neg K_p$) I do not know that p.

So, in virtually every beta-pattern case, B makes A aware of the fact that he does not know that p. Because of this, virtually every beta-pattern case is a case in which B gives A a conclusive epistemic reason to stop actively believing that p.

In what follows, we will refer to this as *the generalised extended LOK explanation* – or, for short, as the *g-extended* LOK explanation. The arguments of sections 1–5 make it clear that this explanation is one that we ought to endorse. In section 6.2, I will use the g-extended LOK explanation to argue that we are all in a position to know the following doctrine:

> (E_{EP}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S is not in a position to know that his belief in P is explained by establishing evidence, then S's belief in P does not constitute knowledge.

Section 6.3 will argue that this doctrine is a form of evidentialism, and section 6.4 will argue that, by appealing to this form of evidentialism, we can generate a problem for ASO_2.

6.2 (E_{EP}) and the g-extended LOK explanation

If the g-extended LOK explanation is true, then the following claim about beta-pattern cases is also true:

> (BPC) In virtually every beta-pattern case, A is aware of the fact that, if he is not in a position to know that (B_p) and ($\neg E_p$) are not both true with respect to him, then ($\neg K_p$) is true with respect to him.

The fact about beta-pattern cases itself stands in need of explanation. For, there are many different kinds of beta-pattern case, involving subjects with very different kinds of background knowledge. How is it that, in virtually *all* of these cases, person A is aware of the truth of the conditional to which (BPC) refers?

One natural way of explaining this fact is to endorse the following claim:

(BPC_2) In virtually every beta-pattern case, it is *a conceptual truth* that, if A is not in a position to know that (B_p) and ($\neg E_p$) are not both true with respect to him, then ($\neg K_p$) is true with respect to him.

It is hard to see how one could explain the truth of (BPC) without endorsing something like this claim. Because of this, we have good reason to believe (BPC_2). But, the truth of (BPC_2) is itself something that stands in need of explanation.

The most natural way of explaining the truth of this claim is to say that there is a broad conceptual truth that entails all of the conceptual truths to which (BPC_2) refers. One simple claim which entails all of the conceptual truths to which (BPC_2) refers is the following (where S's belief that p is *explained by establishing evidence* iff there is some q such that (i) (S believes that p) because q, and (ii) the fact that q establishes that p):

(E) For any subject, S, and proposition P: if S believes P, and S is not in a position to know that his belief in P is explained by establishing evidence, then S does not know P.

There are a few propositions for which (E) may not hold. These propositions are all simple necessary truths, like the proposition that *everything is self-identical*, and the proposition that $2 + 2 = 4$. Although it is clear that we know these propositions, it is by no means clear that we believe them in the way that (E) requires. Because of this, it seems best to explain the truth of (BPC_2) by appealing to the following qualified version of (E):

(E_{EP}) With few exceptions: every proposition P, is such that, for any subject, S: if S believes P, and S is not in a position to know that his belief in P is explained by establishing evidence, then S does not know P.

It is hard to see how we could explain the truth of (BPC_2) without saying that (E_{EP}) is a conceptual truth. Because of this, we have good reason to believe that (E_{EP}) is a conceptual truth. But, if (E_{EP}) is a conceptual truth, then, surely, we are all in a position to *know* (E_{EP}). In the next subsection, I will argue that (E_{EP}) is a form of evidentialism, and in section 6.4, I will use this form of evidentialism to generate a problem for ASO_2.

6.3 From (E_1) to (E_{EP})

A doctrine, D, is a *form of evidentialism* iff D is either identical to the following doctrine, or can be generated by making an admissible amendment to this doctrine:

(E_1) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then S's belief in P is epistemically irrational.

In what follows, we will show that there is an admissible amendment to (E_1) which generates (E_{EP}). By showing this, we will show that (E_{EP}) is a form of evidentialism.

One way of amending (E_1) is to amend the doxastic defect to which it refers. Amendments of this kind are admissible iff the defect to which they appeal is suitable – where a doxastic defect is said to be *suitable* iff it satisfies the following constraint:

(SDD) For any subject, S, and proposition P: if S believes P, and S's belief in P has defect D, then the fact that S's belief has this defect is a conclusive epistemic reason for S to abandon her belief in P.

The arguments of section 5 of chapter two show that the defect of *not constituting knowledge* satisfies (SDD). By showing this, they show that the following doctrine is a form of evidentialism:

(E_2) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S does not have enough evidence for P, then

S's belief in P *does not constitute knowledge.*

Another way of amending (E_1) is to amend the evidential requirement that it imposes on our beliefs. One admissible amendment of this kind is the amendment which generates the following doctrine:

(E_1') With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S is *not in a position to know that his belief in P is based on sufficient evidence*, then S's belief in P is epistemically irrational.

Any amendment to (E_1) which can be brought about by making a series of admissible amendments to this doctrine is itself admissible. So, by combining the amendment which generates (E_2) with the amendment that generates (E_1'), we can generate the following form of evidentialism:

(E_3) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S is not in a position to know that his belief in P is based on sufficient evidence, then S's belief in P does not constitute knowledge.

A third admissible way of amending (E_1) is to replace its constituent terms with terms that can reasonably be believed to be their synonyms. Any amendment to (E_1) which can be brought about by making a series of admissible amendments to this doctrine is itself admissible; so, any amendment to (E_1) which can be brought about by first of all replacing (E_1) with (E_1'), and then replacing terms in (E_1') with terms that can reasonably be believed to be their synonyms, is an admissible amendment to (E_1). It is reasonable to believe that the following schematic sentence:

(BE) S's belief in P is based on sufficient evidence.

is synonymous with this sentence:

(EEE) S's belief in P is explained by establishing evidence.

Because it is reasonable to believe this, it is possible to generate the following doctrine by making an admissible amendment to (E_1):

(E_{EP}) With few exceptions: every proposition P is such that, for any subject S: if S believes P, and S is not in a position to know that his belief in P is explained by establishing evidence, then S's belief in P does not constitute knowledge.

Consequently, we can infer that (E_{EP}) is a form of evidentialism. In what follows, we will refer to it as *epistemic evidentialism.*

6.4 Epistemic evidentialism and the second anti-sceptical objection

In the last two subsections, we have shown that we are all in a position to know a certain form of evidentialism, which we are calling *epistemic evidentialism.* In this section, we will use the results of these subsections to generate a problem for the second anti-sceptical objection. The problem that we will generate is a problem for the second stage of the objection, which is an argument for the following claim:

(2') If explanatory evidentialism does not discredit all deity beliefs, then there is no form of evidentialism that discredits all deity beliefs.

In what follows, we will use the results of sections 6.2 and 6.3 to argue that claim (2') is false.

The first step in our attack on (2') is to argue that, *if explanatory evidentialism discredits all deity beliefs, then epistemic evidentialism also discredits all deity beliefs.* To see that this claim is true, note first of all that a form of evidentialism, E, discredits all deity beliefs iff it is the case that, for every subject, S, who has deity beliefs, the following claims all hold:

(D1) S is in a position to know E, and is in a position to know that her deity beliefs are not exceptions to it.

(D2) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that E imposes, then we can make S aware of this fact.

(D3) (D1) and (D2) are true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way. They are also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

Next, note that, if explanatory evidentialism discredits all deity beliefs, then the following pair of claims hold:

($D2_{EE}$) If S is not aware of the fact that her deity beliefs are not explained by evidence that establishes their truth, then we can make S aware of this fact.

($D3_{EE}$) ($D2_{EE}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way. It is also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

Thirdly: note that, if ($D2_{EE}$) and ($D3_{EE}$) hold, then the following claims also hold (since it is common knowledge that one cannot be in a position to know a proposition that is false):

($D2_{EEP}$) If S is not aware of the fact that she is not in a position to know that *her deity beliefs are explained by evidence that establishes their truth*, then we can make S aware of this fact.

($D3_{EEP}$) ($D2_{EEP}$) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way. It is also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

Finally: note that, if ($D2_{EEP}$) and ($D3_{EEP}$)are true, then epistemic evidentialism discredits all deity beliefs (since, as section 6.2 argued, we are all in a position to know epistemic evidentialism).

The above reasoning shows that, if explanatory evidentialism discredits all deity beliefs, then epistemic evidentialism also discredits such beliefs. It can also be used to show that the reverse implication does not hold. For, it seems clear that ($D2_{EEP}$) does *not* imply ($D2_{EE}$). Consequently, we can conclude that epistemic evidentialism may discredit all deity beliefs even if explanatory evidentialism does not.

Stage two of the second anti-sceptical objection argues that, if explanatory evidentialism does not discredit all deity beliefs, then there is no form of evidentialism that discredits all deity beliefs. We are now in a position to see that this stage of the second anti-sceptical objection fails. The failure of this stage of the second anti-sceptical objection poses a problem for advocates of this objection. In chapter five, I will show that we can solve this problem by constructing another version of the ASO.

Chapter Five

The third anti-sceptical objection

In chapter four, I argued that advocates of the second anti-sceptical objection face a serious problem. In this chapter, I will construct a new version of the anti-sceptical objection, which is invulnerable to this problem. Before I can start to construct this new objection, I need to argue for the following claim:

(M2) If there is a subject, S, who has a self-supporting set of metaphysical deity beliefs, then epistemic evidentialism does not discredit all deity beliefs.

Section 1 clarifies this claim, sections 2–4 argue for it, and section 5 uses it to construct an improved version of the anti-sceptical objection, which I call *the third anti-sceptical objection.*

1. Clarifying (M2)

Before we can defend claim (M2), we must clarify the terminology that it employs. In particular: we need to say something about what it is for a set of metaphysical deity beliefs (henceforth: MDBs) to be *self-supporting.* When I say that a subject, S, has a self-supporting set of MDBs, what I mean is that S has a set of MDBs, M_s, for which the following claim holds:

(SS) Each MDB, M, that belongs to M_s is such that, if S were to reflect about whether M is explained by sensitive evidence[1], then S would properly infer, from members of M_s and knowledge-constituting beliefs, that M *is* explained by sensitive evidence.

[1] S's belief that p is *explained by sensitive evidence* iff there is some q such that (i) (S believes that p) because q, and (ii) if it hadn't been the case that p, then it wouldn't have been the case that q.

Consequently, we can rewrite (M2) in the following way:

(M2') If there is a subject, S, who has a set of MDBs, M_s, for which (SS) holds, then epistemic evidentialism does not discredit all deity beliefs.

In the next three sections, I will give a three-step argument for (M2'). Section 2 will argue that, if the antecedent of (M2') is true, then, under ordinary standards for establishment, the following claim is also true:

($M2_{ae}$') There is a subject, S, who has a set of MDBs, M_s, for which the following claim holds:

(SS_e) Each MDB, M, that belongs to M_s is such that, if S were to reflect about whether M is explained by *establishing* evidence, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that M *is* explained by establishing evidence.

Section 3 will argue that ($M2_{ae}$') implies another, more complex claim which we will call ($M2_{ae}$++'). And, section 4 will argue that, if ($M2_{ae}$++') is true under ordinary standards for establishment, then the consequent of (M2') is true. If the arguments of section 2–4 all succeed, then they will show that the antecedent of (M2') implies the consequent of (M2'). And, if they show this, then they will show that (M2') is true.

2. From ($M2_a$') to ($M2_{ae}$')

In section 2 of chapter 3, we argued that, under ordinary standards for establishment, instances of the following claim:

(C) If it had not been the case that p, then it would not have been the case that q.

obviously imply corresponding instances of this claim:

(E) The fact that q establishes that p.

By appealing to this claim about obvious implication, we can now show that, if the following claim is true:

($M2_a$') There is a subject, S, who has a set of MDBs, M_s, for which the following claim holds:

(SS) Each MDB, M, that belongs to M_s is such that, if S were to reflect about whether M is explained by *sensitive* evidence, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that M *is* explained by sensitive evidence.

then, under ordinary standards for establishment, the following claim is also true:

($M2_{ae}$') There is a subject, S, who has a set of MDBs, M_s, for which the following claim holds:

(SS_e) Each MDB, M, that belongs to M_s is such that, if S were to reflect about whether M is explained by *establishing* evidence, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that M *is* explained by establishing evidence.

To see this, note first of all that to reflect about whether S's belief that p is explained by sensitive evidence is to reflect about whether there is a true instance of the following claim:

(ESE) (S believes that p) because q, and if it hadn't been the case that p, then it wouldn't have been the case that q.

Next, note that to reflect about whether S's belief that p is explained by establishing evidence is to reflect about whether there is a true instance of this claim:

(EEE) (S believes that p) because q, and the fact that q establishes that p.

Finally: note that, if, under ordinary standards for establishment, instances of (C) obviously imply corresponding instances of (E), then, under ordinary standards for establishment, instances of (ESE) obviously imply corresponding instances of (EEE). From these three points, it follows that, if ($M2_a$') is true, then, under ordinary standards for establishment, ($M2_{ae}$') is also true.

In the next section, we will argue that ($M2_{ae}$') implies another, more complex claim, which we will call ($M2_{ae}$++'). During this section, and following sections, we will use the term 'is EEE' to abbreviate the term 'is explained by establishing evidence'.

3. From ($M2_{ae}$') to ($M2_{ae}$++')

Suppose that, as ($M2_{ae}$') claims, there is a subject, S, who has a set of MDBs, M_s, for which the following claim holds:

(SS_e) Each MDB, M, that belongs to M_s is such that, if S were to reflect about whether M is EEE, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that M *is* EEE.

And, suppose that this subject, S, has just properly inferred a belief, B, from members of M_s and knowledge-constituting beliefs, and is now reflecting about whether B is EEE. Since S has just inferred B from members of M_s and knowledge-constituting beliefs, it seems likely that, by introspection, S will come to know that an instance of the following claim is true (where 'B_1', 'B_2' etc. refer either to members of M_s or to knowledge-constituting beliefs):

(i) B has been properly inferred from B_1, B_2, B_3... and B_n.

And, once S has come to know that an instance of this claim is true, it seems likely that S will reflect on whether the corresponding instance of this claim is true:

(ii) B_1 is EEE, and B_2 is EEE...and B_n is EEE.

Because $B_1 - B_n$ are either members of M_s, or knowledge-constituting beliefs, S is likely to conclude that the corresponding instance of (ii) *is* true. For, (SS_e) entails that, when S reflects about whether members of M_s are EEE, he will properly infer, from members of M_s and knowledge constituting beliefs, that they are EEE. And, epistemic evidentialism commits us to saying that, when S reflects about whether his knowledge-constituting beliefs are EEE, he will properly infer, from other knowledge-constituting beliefs, that they are EEE. It seems likely that, once S has come to believe corresponding instances of (i) and (ii), he will properly infer, from these instances of (i) and (ii), that B is EEE. Consequently, it seems likely that, if there is a subject, S, who has a set of MDBs, M_s, for which (SS_e) holds, then the following claim also holds with respect to S:

(I) If S was properly to infer a belief, B, from members of M_s and knowledge-constituting beliefs, and S was then to reflect about whether B is EEE, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that B is EEE.

When (SS_e) is conjoined with (I), it entails the following claim (where 'EM' refers to S's belief that M is EEE, 'EEM' to S's belief that EM is EEE, and so on):

(SS_e+) Each MDB, M, that belongs to M_s is such that the following infinite series of claims is true:

(R_1) If S was to reflect about whether M is EEE, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that M is EEE.

(R_2) If, after reflecting about whether M is EEE, S was to reflect about whether EM is EEE, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that EM is EEE.

(R_3) If, after reflecting about whether M and EM are EEE, S was to reflect about whether EEM is EEE, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that EEM is EEE.

...etc.

And, when this claim is conjoined with the following plausible claims about proper inference:

(PI_1) If a belief, B, is properly inferred from true beliefs, then B will itself be true.

(PI_2) If a belief, B, is properly inferred from beliefs that are EEE, then B will itself be EEE.

it entails the following claim:

(SS_e++) Each MDB, M, that belongs to M_s is such that the following infinite series of claims is true:

(R_{1e}) If S were to reflect about whether M is EEE, then S would come to believe that M is EEE, and would thereby make it the case that EM is EEE.

(R_{2e}) If, after reflecting about whether M is EEE, S was to reflect about whether EM is EEE, then S would come to believe that EM is EEE, and would thereby make it the case that EEM is EEE.

(R_{3e}) If, after reflecting about whether M and EM are EEE, S was to reflect about whether EEM is EEE, then S would come to believe that EEM is EEE, and would thereby make it the case that EEEM is EEE.

...etc.

Consequently, we can infer that, if there is a subject, S, who has a set of MDBs, M_s, for which (SS_e) holds, then (SS_e++) also holds with respect to M_s. And, from this, it follows that, if ($M2_{ae}$') is true, then the following claim is also true:

($M2_{ae}$++') There is a subject, S, who has a set of MDBs, M_s, for which (SS_e++) holds.

In the last section, we showed that, if the antecedent of (M2') is true, then, under ordinary standards for establishment, ($M2_{ae}$') is also true. In this section, we have shown that, if ($M2_{ae}$') is true, then ($M2_{ae}$++') is also true. Consequently, we can now conclude that, if the antecedent of (M2') is true, then, under ordinary standards for establishment, ($M2_{ae}$++') is true. In the next section, we will argue that, if ($M2_{ae}$++') is true under ordinary standards for establishment, then the following claim is true:

($M2_c$') Epistemic evidentialism does not discredit all deity beliefs.

By showing this, we will show that the antecedent of (M2') implies the consequent of (M2'), and will thus show that (M2') is true.

4. From ($M2_{ae}$++') to ($M2_c$')

If, under ordinary standards for establishment, ($M2_{ae}$++') is true, then ($M2_c$') is also true. To see this, we need to recap, briefly, on the definition of discrediting that we gave in chapter one. According to this definition, a form of evidentialism, E, discredits all deity beliefs iff it is the case that, for every subject, S, who has deity beliefs, the following three claims are all true:

(D1) S is in a position to know E, and is in a position to know that her deity beliefs are not exceptions to it.

(D2) If S is not aware of the fact that her deity beliefs do not satisfy the evidential requirement that E imposes, then we can make

S aware of this fact.

(D3) (D2) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether S's epistemic deity beliefs can be rebutted in a non-evidential way. It is also true regardless of whether global scepticism is true, and regardless of whether the standards for epistemic evaluation have been raised.

From this definition, we can infer that, if epistemic evidentialism discredits all deity beliefs, then, for every subject, S, who has deity beliefs, the following pair of claims is true:

(C_1) If S is not aware of the fact that none of her deity beliefs satisfy the evidential requirement that epistemic evidentialism imposes, then we can make S aware of this fact.

(C_2) (C_1) is true regardless of whether S's metaphysical deity beliefs are true, and regardless of whether the standards for establishment have been raised.

If (C_1) and (C_2) are true, for every subject S, then, for every S, the following claim is true, and is true under ordinary standards for establishment:

(C_3) If S's MDBs are all true, and S is not aware of the fact that, for each deity belief, D, that she holds, she is not in a position to know that D is EEE, then we can make S aware of this fact.

But, if the following claim is true, under ordinary standards for establishment:

($M2_{ae}$++') There is a subject, S, who has a set of MDBs, M_s, for which (SS_e++) holds.

then there is some subject for whom (C_3) is *not* true. To see this, note first of all that, if some subject, S, is in a position to know that one of her MDBs is EEE, then the following infinite series of claims is true:

(R_{1e}) If S were to reflect about whether M is EEE, then S would come to believe that M is EEE, and would make it the case that EM is EEE.

(R_{2e}) If, after reflecting about whether M is EEE, S was to reflect about whether EM is EEE, then S would come to believe that EM is EEE, and would make it the case that EEM is EEE.

(R_{3e}) If, after reflecting about whether M and EM are EEE, S was to reflect about whether EEM is EEE, then S would come to believe that EEM is EEE, and would make it the case that EEEM is EEE.

...etc.

Next, note that the most promising way of making a subject, S, aware of the fact that she is not in a position to know that an MDB, M, is EEE, is to make S aware of the fact that one of the claims in the series (R_{1e}), (R_{2e}), (R_{3e})... is false. These two points strongly suggest that, if, for every subject S, (C_3) is true under ordinary standards for establishment, then, for every S, the following claim is true under ordinary standards for establishment:

(C_3') If S's MDBs are all true, then, for each MDB, M, that she holds, one of the claims in the series (R_{1e}), (R_{2e}), (R_{3e})... is false.

But, if ($M2_{ae}$++') is true, under ordinary standards for establishment, then there is some subject for whom (C_3') is not true, under ordinary standards for establishment.

The upshot of this section is that, if, under ordinary standards for establishment, ($M2_{ae}$++') is true, then the following claim is also true:

($M2_c$') Epistemic evidentialism does not discredit all deity beliefs.

When the conclusion of this section is conjoined with the conclusions of the last two sections, it entails that the following claim is true:

(M2) If there is a subject, S, who has a self-supporting set of metaphysical deity beliefs, then epistemic evidentialism does not discredit all deity beliefs.

In the next section, I will show that, by appealing to this claim, we can construct an improved version of the second anti-sceptical objection.

5. The third anti-sceptical objection

In this section, I will describe an improved version of the second anti-sceptical objection (henceforth: ASO_2). I will refer to this objection as the third anti-sceptical objection – or, for short, as ASO_3. ASO_3 is a four-stage argument which aims to give people a conclusive epistemic reason to stop actively believing EBG. In what follows, I will describe each stage of this argument.

The first stage of ASO_3 is the argument that we gave in chapter one for the following claim:

(1) If EBG is true, then there is a form of evidentialism that discredits all deity beliefs.

Nothing that we have said in chapters 2–4 threatens the argument that we gave for (1); so, we can safely retain this argument as a part of ASO_3.

The second stage of ASO_3 is an argument for a qualified version of this claim:

(2) There is no form of evidentialism that discredits all deity beliefs.

We can generate this argument by focusing on the second stage of the original ASO. According to this stage of the objection, virtually every form of evidentialism either leads to scepticism, or is too weak to discredit all deity beliefs. In chapter two, we showed that explanatory evidentialism does not succumb to this stage of the objection; and, in chapter four, we showed that the same is true with respect to epistemic evidentialism. But, the arguments of chapter four strongly suggest that there are no other forms of evidentialism that are invulnerable to this stage of the objection. Because of this, the second stage of the original ASO gives us reason to believe the following, qualified version of claim (2):

(2") If neither explanatory evidentialism nor epistemic evidentialism discredit all deity beliefs, then there is no form of evidentialism that discredits all deity beliefs.

And since any deity belief that is discredited by explanatory evidentialism is also discredited by epistemic evidentialism, we can simplify this claim in the following way:

(2"') If epistemic evidentialism does not discredit all deity beliefs, then there is no form of evidentialism that discredits all deity beliefs.

The third stage of ASO_3 is a three-part argument for the following claim:

(3") If there is some subject, S, who has a self-supporting set of MDBs, then EBG is not true.

The first part of this argument points out that claims (1) and (2"') jointly entail the following claim:

(3a") If epistemic evidentialism does not discredit all deity beliefs, then EBG is not true.

The second part uses the arguments of sections 1–3 to defend this claim:

(M2) If there is a subject, S, who has a self-supporting set of metaphysical deity beliefs, then epistemic evidentialism does not discredit all deity beliefs.

And, the third part points out that (3a") and (M2) jointly entail (3").

The fourth and final stage of ASO_3 is an argument for this claim:

(4") Under ordinary standards for knowledge, EBG is not known to be true.

To state this argument, we need to focus on the following claim:

(D*) There is a community of people who all believe in a certain deity (who we will call *D**). Many of the things that the members of this community believe about D* are believed by them because they are written in a certain book (which we will call *the D*-book*). According to this community of D*-believers, the D*-book is, in a certain sense, a book that was written by D* himself. The D*-believers also think that D* would never write a book which said something untrue, and they think this for the sense of 'write' which they use when they say that *the D*-book was written by D**. The claim that *the D* book was written by D**, and the claim that *D* would never write something untrue* are both written in the D*-book. And, both claims are believed by the D*-believers *because* they are written in this book.

The fourth stage of ASO_3 argues for claim (4") by arguing for the following pair of claims:

(4a") If (D*) is true, then EBG is not true.

(4b") If (D*) is not true, then, under ordinary standards for knowledge, EBG is not known to be true.

In the next four subsections, I will state the fourth stage of ASO_3 by stating the arguments from which it is composed.

5.1 *Arguing for (4a")*

If (D*) is true, then there are people who have MDBs with the following content:

(MDB_1) The D*-book is a book that was, in a certain sense, written by D* himself.

(MDB_2) D* would never write a book which said something untrue.

Let us focus on one of these people, who we can call S. In what follows, I will argue that S has a set of MDBs which is self-supporting – i.e. a set of MDBs, M_s, for which the following claim is true:

(SS) Each MDB, M, that belongs to M_s is such that, if S were to reflect about whether M is explained by sensitive evidence, then S would properly infer, from members of M_s and knowledge-constituting beliefs, that M *is* explained by sensitive evidence.

Consider the set of MDBs that contains just S's belief in (MDB_1), and S's belief in (MDB_2). This set of beliefs, which we will call M_s*, is self-supporting. To see that it is self-supporting, consider S's belief in (MDB_1). (D*) tells us that S believes (MDB_1) because (MDB_1) is written in the D-book. Consequently, it is natural to say that, if S were to reflect about whether his belief in (MDB_1) is explained by sensitive evidence, then S would, by introspection, come to know that the following claim is true with respect to him:

(B_{sm}) I believe (MDB_1) because (MDB_1) is written in the D*-book.

It is also natural to say that, under these circumstances, S would properly infer, from (MDB_1) and (MDB_2), that the following claim is true:

(C_{sm}) If (MDB_1) had not been true, then (MDB_1) wouldn't have been written in the D*-book.

For, it is clear that this claim can properly be inferred from (MDB_1) and (MDB_2). And, since S is reflecting about whether his belief in (MDB_1) is explained by sensitive evidence, it seems likely that he would notice that this inference is available to him.

After S had come to believe (B_{sm}) and (C_{sm}), he would properly infer, from these two claims, that his belief in MDB_1 is explained by sensitive evidence. So, if S were to reflect about whether his belief in (MDB_1) is explained by sensitive evidence, then S would properly infer, from members of M_s* and knowledge-constituting beliefs, that his belief in (MDB_1) *is* explained by sensitive evidence. The above reasoning can be duplicated with respect to S's belief in (MDB_2). So, since M_s* contains just S's belief in (MDB_1) and S's belief in (MDB_2), we can conclude that M_s* is self-supporting.

The above reasoning shows that, if (D*) is true, then there are subjects who make the antecedent of the following claim true:

> (3") If there is some subject, S, who has a self-supporting set of MDBs, then EBG is not true.

The third stage of ASO_3 shows that claim (3") is true; so, we can now conclude that, if (D*) is true, then the consequent of (3") is true; or, equivalently, that the following claim is true:

> (4a") If (D*) is true, then EBG is not true.

5.2 *Arguing for (4b")*

Before arguing for (4b"), we must argue for another claim. This claim can be stated in the following way:

> (EP_1*) If D* is not true, then it could easily have been true.

In the next section, I will argue for this claim. Section 5.4 will then use this claim, and the following claim (which was defended in section 4.4 of chapter 3), to argue for (4b'):

> (EP_2) If it could easily have been the case that p, then, under ordinary standards for establishment, instances of the following claim:
>
> > ($\neg C$) If it had not been the case that p, then it would still have been the case that q.
>
> imply corresponding instances of this claim:
>
> > ($\neg E$) The fact that q does not establish that p.

5.3 Arguing for (EP_1)

There are certain actual communities which closely resemble the community described in claim (D*). One such is the community of Christians. Many of the things that Christians believe about God are such that Christians believe them because (or at least: partly because) the Bible says that they are true. And, many Christians seem to think (a) that the Bible was, in a certain sense, written by God, (b) that God would never write a book which said something untrue, and (c) that claims (a) and (b) can be properly inferred from claims that are written in the Bible.

It may turn out that the community of Christians does not conform exactly to our description of the D*-believers. And, it may turn out that there is no other community which conforms exactly to this description. But, it seems clear that there are a number of communities that come *close* to conforming to this description. And, for this reason, it is natural to endorse the following claim:

> (EP_1) If (D*) is not true, then it could easily have been true.

5.4 From (EP_1) to (4b')

Now that we have defended (EP_1), we can start to argue for (4b'). The first step in our argument is to point out that, if (D*) is not true, then the following, counterfactual claim will surely hold:

(i) Even if (D*) was true, advocates of EBG would believe EBG, and their reasons for believing it would be exactly the same as they actually are.

This counterfactual claim implies the following claim (where 'R' refers to the reason for which advocates of EBG believe EBG, whatever that reason is):

(i') If (D*) was true, then advocates of EBG would still believe EBG, and would believe it because R.

The next step in our argument is to point out that (i') entails this claim:

(ii) If (D*) was true, then it would be the case that R.

And, the third step is to point out that, when (ii) is conjoined with (EP_1) and (EP_2), it entails the following claim:

(iii) Under ordinary standards for establishment, the fact that R does not establish that (D*) is not true.

The fourth step in our argument points out that, since we can infer the falsity of (D*) from the truth of EBG, the following claim is true:

(iv) If R doesn't establish that (D*) is not true, then R does not establish that EBG is true.

From (iii) and (iv), it follows that, under ordinary standards for establishment, R does not establish that EBG is true. So, when (iii) and (iv) are conjoined with explanatory evidentialism, they entail that, under ordinary standards for knowledge, EBG is not known to be true. Because of this, we can conclude that the following claim is true:

(4b'') If (D*) is not true, then, under ordinary standards for knowledge, EBG is not known to be true.

5.5 Summary

If the arguments of the last five subsections succeed, then they show that the following claims are both true:

(4a'') If (D*) is true, then EBG is not true.

(4b'') If (D*) is not true, then, under ordinary standards for knowledge, EBG is not known to be true.

From this pair of claims, we can infer this claim, which ASO_3 aims to establish:

(4'') Under ordinary standards for knowledge, EBG is not known to be true.

If ASO_3 does establish this claim, then it can be used to give people a conclusive reason to stop actively believing EBG. Consequently, it is worth investigating whether ASO_3 establishes this claim. I cannot embark on this investigation here. But, I suspect that it will yield positive results for advocates of ASO_3.

6. Summary of thesis

In this thesis, I have developed three anti-sceptical objections to the following claim:

(EBG) It is possible to undermine belief in God by arguing that the evidence for this belief is insufficient.

The first of these objections was outlined in chapter one; the second was outlined in chapter three; and the third has been outlined in this chapter. It is clear that the third anti-sceptical objection is superior to its predecessors, but it is not yet clear that this objection succeeds. In future work, I hope to show that it does succeed. If this work is successful, then it will show that all of us have, or can be given, a conclusive epistemic reason to refrain from endorsing EBG.

References

Blanshard, B. 1974. *Reason and Belief.* London: Allen and Unwin.

Clifford, W.K. 1879. The ethics of belief. In his *Lectures and Essays*. London: Macmillan.

Cohen, L.J. 1992. *An Essay on Belief and Acceptance*. Oxford: Oxford University Press.

Cohen, S. 1988. How to be a fallibilist. In J. Tomberlin (ed.), *Philosophical Perspectives 2*. Atascadero, CA: Ridgeview.

Cohen, S. 1999. Contextualism, skepticism and reasons. In J. Tomberlin (ed.), *Philosophical Perspectives 13*. Oxford: Blackwell.

de Sousa, R. 1971. How to give a piece of your mind: or, the logic of belief and assent. *Review of Metaphysics*, 25(1): 52-79.

DeRose, K. 1991. Epistemic possibilities. *Philosophical Review*, 100: 581-605.

DeRose, K. 1995. Solving the skeptical problem. *Philosophical Review* 104: 1-52.

DeRose, K. 1996. Knowledge, assertion and lotteries. *Australiasian Journal of Philosophy,* 74: 568-80.

DeRose, K. 2000. Ought we to follow our evidence? *Philosophy and Phenomenological Research* 56(1): 193-7.

Dretske, F. 1970. Epistemic operators. *Journal of Philosophy* 67: 1007-23.

Dretske, F. 1971. Conclusive reasons. *Australasian Journal of Philosophy*, 49: 1-22.

Feldman and Conee. 1985. Evidentialism. *Philosophical Studies* 48: 15-34.

Flew, A. 1976. *The Presumption of Atheism*. London: Pemberton.

Frankish, K. 1998. Natural language and virtual belief. In P. Carruthers and J. Boucher (eds.), *Language and Thought*. Cambridge University Press.

Grice, H.P. 1975. Logic and conversation. In D. Davidson and G. Harman (eds.), *The Logic of Grammar*. Encino: Dickensen.

Harman, G. 1968. Knowledge, inference and explanation. *American Philosophical Quarterly* 5: 164-73.

Hambourger, R. 1987. Justified assertion and the relativity of knowledge. *Philosophical Studies*, 51: 241-69.

James, W. 1896. The will to believe. In *The Will to Believe and Other Essays in Popular Philosophy*. London: Longmans.

Mackie, J.L. 1982. *The Miracle of Theism*. Oxford: Oxford University Press.

Moore, G.E. 1962. *Commonplace Book: 1919-1953*. London: Allen & Unwin.

Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.

Plantinga, A. 1981. Is belief in God properly basic? *Nous* 15: 41-51

Plantinga, A. 1983. Reason and belief in God. In A. Plantinga and N. Wolterstorff (eds.), *Faith and Rationality*. Notre Dame: University of Notre Dame Press.

Plantinga, A. 1985. Reply to Alston. In J. Tomberlin and P. van Inwagen (eds.), *Alvin Plantinga*. Dordrecht: Kluwer Academic Press.

Plantinga, A. 1986. Coherentism and the evidentialist objection. In R.Audi and W.J. Wainwright (eds.), *Rationality, Religious Belief and Moral Commitment*. Ithaca: Cornell University Press.

Plantinga, A. 2000. *Warranted Christian Belief*. New York: Oxford University Press.

Russell, B. 1957. Why I am not a Christian. In his *Why I am not a Christian*. New York: Simon and Schuster.

Slote, M. A. 1979. Assertion and belief. In J. Dancy (ed.), *Papers on Language and Logic*. Keele: Keele University Press.

Sorensen, R.A. 1988. *Blindspots*. New York: Oxford University Press.

Sorensen, R.A. 2000. Moore's problem with iterated belief. *Philosophical Quarterly* 50: 28-43.

Swinburne 1979. *The Existence of God*. Oxford: Oxford University Press.

Unger, P. 1975. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.

van Inwagen, P. 1997. It is wrong, everywhere, always and for anyone to believe anything on insufficient evidence. In his *'The Possibility of Resurrection' and Other Essays in Christian Apologetics*. Westview Press.

Williamson, T. 1992. Inexact knowledge. *Mind*, 101: 217-42.

Williamson, T. 1996. Knowing and asserting. *Philosophical Review* 105: 489-523.

Williamson, T. 1997. Knowledge as Evidence. *Mind*, 106: 717-41.

Williamson, T. 2000. *Knowledge and its Limits*. New York: Oxford University Press.

Wykstra, S. 1989. Towards a sensible evidentialism: On the Notion of 'Needing Evidence'. In W.L. Rowe and W.J. Wainwright (eds.), *Philosophy of Religion: Selected Readings*. New York: Harcourt Brace Jovanovich.