

SCPTICAL THEISM AND EVIDENTIAL ARGUMENTS FROM EVIL

Michael J. Almeida and Graham Oppy

Sceptical theists—e.g., William Alston and Michael Bergmann—have claimed that considerations concerning human cognitive limitations are alone sufficient to undermine evidential arguments from evil. We argue that, if the considerations deployed by sceptical theists are sufficient to undermine evidential arguments from evil, then those considerations are also sufficient to undermine inferences that play a crucial role in ordinary moral reasoning. If cogent, our argument suffices to discredit sceptical theist responses to evidential arguments from evil.

In recent times, a number of philosophers have championed ‘sceptical theist’ responses to evidential arguments from evil [Alston 1991, 1996; Bergmann 2001; Fitzpatrick 1981; Howard-Snyder 1996a, 1996c; van Inwagen 1991; Plantinga 1979, 1988; Wykstra 1984, 1996]. The core idea behind these responses to evidential arguments from evil is that considerations of human cognitive limitations are *alone* sufficient to undermine those arguments. This core idea is developed in different ways. Some ‘sceptical theists’ [Bergmann 2001; Howard-Snyder 1996a, 1996c] claim that consideration of human cognitive limitations in the realm of value are alone sufficient to undermine evidential arguments from evil. Other ‘sceptical theists’ [Alston 1991, 1996; van Inwagen 1991; Plantinga 1979, 1988; Wykstra 1984, 1996] claim that consideration of human cognitive limitations in various spheres including the realm of value are alone sufficient to undermine evidential arguments from evil. Our response to these ‘sceptical theists’ is in two parts. First, we argue—against [Bergmann 2001, et al.]—that it isn’t true that considerations of human cognitive limitations in the realm of value are *alone* sufficient to undermine evidential arguments from evil. Second, we argue against [Alston 1991, 1996, et al.]—that it isn’t true that considerations of human cognitive limitations in various spheres including the realm of values are *alone* sufficient to undermine evidential arguments from evil.

The structure of the paper is as follows. In the first section, we provide what we take to be a more or less canonical formulation of an evidential argument from evil. This argument differs in some ways from arguments which have hitherto been presented in the literature, but nothing turns on these differences: the subsequent argument could be developed equally well in connections with William Rowe’s more familiar formulations of evidential arguments from evil.¹ What is crucial is that the kind of evidential argument which we consider—and which is the target of ‘sceptical theist’ critique—contains a key inference

¹ See Rowe [1979] and many subsequent publications. In more recent publications, Rowe has moved away from arguments that involve the kind of inference that will be the main focus of our discussion in this paper. However, there is a fairly straightforward variant of the sceptical theist response which can be made to Rowe’s new arguments, and exactly the same kinds of considerations will arise. So there is no loss of generality in focussing on the particular evidential argument from evil that we set out below.

which moves from a premise of the form ‘We have found no reasons why...’ to a conclusion of the form ‘There are no reasons why...’. The burden of the ‘sceptical theist’ critique of the argument is to attack this inference by appeal to nothing more than certain (alleged) human cognitive limitations; we shall argue in reply, that the inference in question cannot be undermined merely by appeal to considerations about human cognitive limitations in the realm of value, or by appeal to considerations about human cognitive limitations in a variety of spheres including the realm of value.

In the second section of the paper, we discuss Stephen Wykstra’s claim that the premise in the key inference does not even weakly support the conclusion of that inference, i.e., the claim that our failure to find reasons of a certain kind does not raise the probability that there are no reasons of that kind *at all*.² We suggest that this claim is massively implausible: it would require very special circumstances in order for a claim of the form ‘We have found no reasons why...’ to fail to provide *any* support for the claim that ‘There are no reasons why...’; and it is not credible to suppose that our evidential argument from evil provides such a case.

In the third section of the paper, we move on to discuss Michael Bergmann’s ‘sceptical theist’ argument—based on a ‘general scepticism about our knowledge of the realm of value’—that the premise in the key inference does not provide substantial support for the conclusion of that inference [Bergmann 2001].³ We claim that, if the inference from ‘We have found no reasons why...’ to ‘There are no reasons why...’ is blocked by the considerations to which Bergmann adverts in the case of evidential arguments from evil, then similar inferences will be blocked in cases of ordinary moral reasoning which we all have reason to endorse. Thus, we claim, Bergmann faces a dilemma: either his ‘general scepticism about our knowledge of the realm of value’ is too benign to save theism from the evidential argument from evil; or else his ‘general scepticism about our knowledge of the realm of value’ is so strong that it threatens to disrupt our ordinary patterns of moral reasoning.

In the fourth and fifth sections of our paper, we turn to consider possible objections to our argument against Bergmann. In particular, we consider the suggestion that our argument fails to pay due attention to the fact that it makes a huge difference that it is not we but rather a perfect being who is the subject of the key inference in our evidential argument from evil; and we also consider the suggestion that we are wrong to think that the moral scepticism which is mandated by the principles that Bergmann endorses is anything other than benign.

In the sixth section of our paper, we consider the prospects for extending our critique of Bergmann [2001] to other kinds of ‘sceptical theist’ responses to evidential argument from evil. In particular, we focus on Alston [1991, 1996]. Unlike Bergmann, Alston does not claim that considerations concerning human cognitive limitations in the sphere of value are alone sufficient to undermine evidential arguments from evil. However, it seems to us that Alston is committed to the claim that considerations concerning human cognitive limitations in a range of spheres including the realm of values are alone sufficient to undermine evidential arguments from evil. But—as we shall go on to

² This claim is made in Wykstra [1984]; it is retracted in Wykstra [1996].

³ Rowe has a reply to Bergmann in the same volume of the journal. However, we shall not discuss this reply, since it takes a quite different tack from the one that we wish to pursue.

argue—if the objection that we develop against Bergmann is good, then it carries over to this view as well. Moreover, if we suppose—as we think we should—that it is constitutive of ‘sceptical theist’ responses to evidential arguments from evil to claim that considerations concerning human cognitive limitations in a range of spheres including the realm of values are alone sufficient to undermine evidential arguments from evil, then our argument carries over to all versions of ‘sceptical theist’ responses to evidential arguments from evil.⁴

Perhaps it is worth emphasising here that we are not setting out to defend the claim that some evidential arguments are successful pieces of atheological argumentation. Rather, what we are aiming to do is to defend these arguments from one kind of theistic counter-attack. We are prepared to allow that, for example, the ‘G.E. Moore shift’ described by Rowe [1979] provides a perfectly satisfactory response to these arguments. (Why shouldn’t the theist presented with Rowe’s argument infer that it is very likely that there are unknown goods that justify God in permitting certain kinds of evils? Such an inference is entirely compatible with the further claim that we know that there are no unknown goods which bear on our own moral evaluations—not least, perhaps, because we have a guarantee from God that this is so.) However, it is important to note that the ‘sceptical theist’ response is distinctive because, were it correct, it would establish that no-one—theist or non-theist—should make the crucial probability judgement that undergirds evidential arguments from evil. The direction of the ‘sceptical theist’ argument is from considerations concerning our cognitive limitations—and nothing but our cognitive limitations—to the conclusion that a certain probability judgement is out of bounds. And it is this that we claim to be able to show is wrong.

I

We begin, then, with the formulation of a more or less canonical version of an evidential argument from evil. The main sub-conclusion of the evidential argument from evil that we shall discuss is that there are evils for which it is true that, were a perfect being to prevent those evils, then the world would be non-arbitrarily improved thereby. It follows from this sub-conclusion that, if there are such evils, then there is no perfect being. It is tempting to call such evils ‘gratuitous’; however, it seems doubtful that this usage is in accord with the standard definition(s) in the literature.⁵ So let us instead call them ‘problematic’.⁶

⁴ Perhaps there is some room for dispute about whether the other writers to whom we have referred conform to this account of ‘sceptical theism’. Our target in this paper is the view that we have just characterized; we’re happy to defer argument about exactly who has defended this kind of view in print to some other occasion. We think that it is clear that this view is defended by Bergmann, Alston, Howard-Snyder, etc.; but we don’t propose to argue the remaining cases here.

⁵ For instance, some people have claimed that God is not *required* to prevent any evil whose existence is logically entailed by some greater overall good. On the assumption that ‘evils’ are evil states of affairs, and ‘goods’ are good states of affairs, the further assumption that there are conjunctive states of affairs which are goods, even though they contain conjuncts which are evils, will pretty quickly lead to the conclusion that almost all evils are such that God is not required to prevent them. While there are ways of avoiding the trivialization of this kind of conception of ‘gratuitous’ evil, it seems to us to be preferable to look for a different way of thinking about the kinds of evils which God would be obliged to prevent.

⁶ It should be noted that we have taken no stand on the question of what is required in order for the prevention of an evil to non-arbitrarily improve the world. Perhaps what is required is, at least, that

Evidential arguments from evil—of the kind which we are investigating here—aim to establish that there are problematic evils, or that there is most reason to think that there are problematic evils, or that there is at least some reason to think that there are problematic evils. (Later, we shall fuss more about the strength of the conclusion which is to be drawn in arguments of this kind; for now, we leave the matter open.)

Before we turn to the formulation of our target argument, we should explain what we mean when we insist that problematic evils are evils that are such that, if they were prevented by the actions of a perfect being, then the world would be *non-arbitrarily* improved.

As a first step towards motivating our account of problematic evils, it should be noted that it would not suffice for the purposes of the argument to claim that there are evils for which it is true that, were a perfect being to prevent those evils, then the world would not thereby be made worse. Suppose that there are two worlds, W_1 and W_2 which are in all respects alike except that W_1 contains evil E_1 and W_2 contains E_2 , where E_1 and E_2 are of equal (dis)value. If W_1 is actual, then it casts no doubt on the claim that there is a perfect being to point out that a perfect being could have prevented the occurrence of E_1 —by actualising world W_2 instead—without thereby making the world worse. (We assume—what some have denied—that, if a perfect being must choose between W_1 and W_2 , then it can do so.)

As a more important step towards motivating our account of problematic evils, it should be noted, too, that it would not suffice for the purposes of the argument to claim that there are evils for which it is true that, were a perfect being to prevent those evils, then the world would be improved. Suppose that there is a sequence of worlds $W_1, W_2, \dots, W_n, \dots$ which contain sequences of evils $\{E_1, E_2, \dots, E_n, \dots\}, \{E_2, E_3, \dots, E_n, \dots\}, \dots, \{E_k, E_{k+1}, \dots, E_n, \dots\}, \dots$, and which are otherwise identical. (So, by hypothesis, W_1 is worse than W_2 , which is worse than W_3 , etc.) Suppose further that any world which contains only a finite number of the E_i is worse than any of the worlds $W_1, W_2, \dots, W_n, \dots$, because a final infinite segment of the sequence of evils $E_1, E_2, \dots, E_n, \dots$, is required in order to ensure some massive ‘outweighing’ good. If W_k is actual, it casts no doubt on the claim that there is a perfect being to point out that a perfect being could have prevented the obtaining of E_k —by actualising world W_{k+1} , say—and thereby made the world better. (Again, we assume—what some may deny—that if a perfect being has to choose one final infinite section of the sequence $E_1, E_2, \dots, E_n, \dots$, then it can do so.)

The point here is that it is possible that a perfect being may need to make an arbitrary choice, from amongst a set of worlds which it can make, if it is to make any world at all. If, for example, for any world that a perfect being can make, there is a better world, then, if the perfect being is to make any world at all, it must arbitrarily choose one which is good enough. But, in these circumstances, it is then no criticism of the activity of the perfect being that it could have made a better world: in the circumstances, the mooted

⁶ *continued*

the ‘net value’ of the world is non-arbitrarily increased; perhaps what is required is, at least, that the ‘net value’ of the world should not decrease while the ‘net evil’ is non-arbitrarily decreased; perhaps what is required is, at least, that the ‘net evil’ in the world is non-arbitrarily decreased sufficiently to justify a corresponding decrease in the ‘net value’ of the world; perhaps what is required is something else. One suggestion we cannot accept is that what is required is that the ‘net value’ of the world is maximised; this suggestion is plainly too demanding. The subsequent discussion does not require that we take any stance on this issue.

improvement of the world is *arbitrary*. The case that we have described in the previous paragraph is simply meant to show that improvements of the world by the prevention of evils *could* be arbitrary in exactly the same sense: there may be evils which are not required for the obtaining of greater goods, or for the non-obtaining of greater evils, and yet which a perfect being does not have reason to prevent. However, we do assume—what some may perhaps deny—that, if the world would be *non-arbitrarily* improved by the prevention of some evil then, a fortiori, a perfect being would prevent that evil; in other words, if there are problematic evils, then there is no perfect being.

With these preliminaries out of the way, we can now turn to our formulation of an evidential argument from evil. Let E be some candidate—i.e., *prima facie*—problematic evil which has occurred, e.g., the rape, beating, and murder by strangulation of a five year old girl; or the prolonged and painful death of a fawn which has been trapped in a forest fire; or the like. We shall use the following as our representative evidential argument from evil:

- (1) We have been unable to find even *pro tanto* reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. (Premise)
- (2) (Therefore) There are not even *pro tanto* reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. (From 1)
- (3) There are at least *pro tanto* reasons why the world would be non-arbitrarily improved if a perfect being prevented E. (Premise)
- (4) (Therefore) There is all-things-considered reason why the world would be non-arbitrarily improved if a perfect being prevented E. (From 2, 3)
- (5) (Therefore) The world would be non-arbitrarily improved if a perfect being prevented E. (From 4)
- (6) (Therefore) There is no perfect being. (From 5)

In support of (1), we note that we can point to no greater good which would be lost if E were prevented by a perfect being; no greater evil which would ensue if E were prevented by a perfect being; nothing which suggests that there must be some events like E—violations of young children—if there are to be greater goods which would otherwise be lost; nothing which suggests that there must be some events like E—violations of young children—if there are not to be greater evils which would otherwise ensue.⁷

In support of (3), we note that we all agree that the world could have been non-arbitrarily improved if one of us had intervened to prevent E, other circumstances permitting. Indeed, we note that we all agree that, other circumstances permitting, we would have a moral obligation to intervene: if we could, without risk to ourselves and others, and without extravagant use of resources, prevent E, then that is what we are required to do. (No doubt, our moral obligation runs further than this; however, all we need is the uncontroversial claim that our obligations run at least this far.) The awfulness of E is enough to establish that a perfect being has at least a *pro tanto* reason to prevent it.

The inference from (4) to (5) is uncontroversial. The inference from (2) and (3) to (4) also looks solid: if there are not even *pro tanto* reasons why the world would not be

⁷ Perhaps the argument in this paragraph is more controversial than we have allowed. For instance, if van Inwagen [1991] were right, then we do have some reason to think that, if the world is to be law-governed, then there must be some events like E. We think that van Inwagen is not right; but we don't propose to try to argue the case here.

non-arbitrarily improved if a perfect being prevented E, and there are *pro tanto* reasons why the world would be non-arbitrarily improved if a perfect being prevented E, then it surely follows that there is all-things-considered reason why the world would be non-arbitrarily improved if a perfect being prevented E.⁸ That leaves the inference from (1) to (2); and here there are many philosophers who will want—and have wanted—to raise objections. In what follows, we shall consider some of the objections that have been made to the inference from (1) to (2).

II

The inference from (1) to (2) moves from ‘It is not the case that we have found reasons of such-and-such a kind’ to ‘There are no reasons of such-and-such a kind’.⁹ Plainly, this is not in general a good deductive inference: it is perfectly possible for claims of the former kind to be true while the corresponding claims of the latter kind are false. However, it does not follow from this claim alone that the inference from (1) to (2) is not a good inference; it may be that there is some other kind of evidential or probabilistic support which (1) lends to (2). Moreover, it seems natural—at least initially—to suppose that this is the case: surely (1) does lend some kind of evidential or probabilistic support to (2). That I have failed to find reasons of such-and-such kind may well be evidence for the claim—may well make it more likely—that there are no reasons of such-and-such kind.¹⁰

We need to distinguish at least two different possible claims here. One claim is that (1) provides *strong evidential* support for (2): given (1), we have substantial reason—perhaps even more reason than not—to believe (2). A weaker claim is that (1) provides *weak evidential* support for (2): given (1), we have more reason to believe (2) than we had before we took (1) into account. We shall first briefly consider views that dispute even the weaker claim, before turning to consideration of arguments against the stronger claim.

A clear example of someone who disputes even the weaker claim is Wykstra [1984]. In his view: ‘Cognizance of suffering... *should not in the least* reduce our confidence that [perfect being theism] is true. When cognizance of suffering does have this effect, it is perhaps because we had not understood [what perfect being theism] proposes for belief in the first place’ [Wykstra 1984: 91, our italics]. Of course, this kind of position is very strong. If cognizance of suffering should not in the least reduce our confidence that perfect being theism is true, then it seems that observations of evils in the world must be completely irrelevant for the question of the assessment of the truth of the claim that there is a perfect

⁸ Note that the argument which we offer here is not the argument given by Russell [1989]: we are not supposing that it follows, from the fact that (7) there is all-things-considered reason why the world would be non-arbitrarily improved if *we* prevented E, that (4) there is all-things-considered reason why the world would be non-arbitrarily improved if *a perfect being* prevented E. Perhaps it might be argued that (7) is at least weak evidence for (4); however, there are obvious difficulties which confront the claim that (7) is strong evidence for (4). In particular, it could well be that, even though the world would be non-arbitrarily improved by *our* preventing E, the world would not be non-arbitrarily improved by *a perfect being's* preventing E.

⁹ In the literature, inferences of this kind are sometimes referred to as ‘noseum inferences’, after Wykstra [1996]. We shall sometimes make use of this label.

¹⁰ To strengthen the case for this claim, we can add that the investigation has been neither careless nor casual: many people have devoted extensive effort to the search for reasons of the kind in question.

being. Suppose, for example, that we were to discover that there are a billion other inhabited planets in our galaxy, and that the trillions of intelligent inhabitants of all those other planets live lives of unrelieved misery. Suppose that we extend our search, and find the same figures reproduced for the millions of other galaxies. Suppose, finally, that we are utterly unable to think of any way in which all of this misery could subserve some greater good. The view in question entails that *none* of this would provide any evidence at all against perfect being theism. It is very tempting to suppose that this is a *reductio* of the view in question.

What would one need to believe in order to defend the claim that (1) does not even provide weak evidential support for (2)? Well, the claim is that F, the failure to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E, provides no support at all for the hypothesis, R, that there is no reason why the world would not be non-arbitrarily improved if a perfect being prevented E. Since, by Bayes's Theorem, $\Pr(R/F) = \Pr(R)$. $\Pr(F/R)/\Pr(F)$, it follows that what needs to be believed is that $\Pr(F/R) \leq \Pr(F)$. (Here, we assume that F weakly supports R exactly if $\Pr(R/F) > \Pr(R)$.) That is, we need to think that it is no less likely that *we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E* than it is that *we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E given that there is no reason why the world would not be non-arbitrarily improved if a perfect being prevented E*. But it is *certain* that, if there is no reason why the world would not be non-arbitrarily improved if a perfect being prevented E, then we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. In other words, the conditional probability which we should assign here is as high as a conditional probability can be. So what is required—in order to defend the position which Wykstra espouses—is that we assign a probability of *no less than one* to the claim that we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. But that is absurd; none of us should assign a probability of no less than one—i.e., of exactly one—to the claim that we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E.¹¹

Wykstra himself identified a different condition that he supposed would need to be satisfied in order for (1) to provide weak evidential support for (2). In his view, (1) cannot provide even weak evidential support for (2) unless it is true that, were there a perfect being that had reasons for not preventing E, we would be able to find those reasons. Moreover, Wykstra also holds that, in fact, it isn't true that were there a perfect being that had reasons for not preventing E, we would be able to find those reasons. Why should we accept this claim? Because we should accept that, were there a perfect being that had reasons for not

¹¹ Note that, even if we were to suppose that the conditional probability in question is merely *very high*—perhaps because we mistakenly supposed that it is relevant to note that we could mistakenly suppose that there are reasons when in fact there are none—it would still be the case that Wykstra's position should be rejected. Under this revised supposition, it would be the case that what needs to be argued for is the claim that we should assign a more than very high probability to the claim that we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. But, even if one thinks that it is certain that there is a perfect being, it is hard to see how one could be justified in supposing that it is so close to certain that we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. What justification is there, for instance, for being so very confident that we shall never construct a theodicy which is able to provide such reasons?

preventing E, we would not even be able to *understand* those reasons.

The discussion in the previous paragraph but one suggests several reasons for dissatisfaction with this argument. As we noted there, what needs to be argued is that it is *very* likely that we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. Even if we claimed that it is *certain* that, were there a perfect being which had reasons for not preventing E, we would not be able to understand those reasons—something which plainly goes far beyond what it is reasonable to claim—it would not follow that it is *very* likely that we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E. If we give a realistic assessment of the likelihood that, were there a perfect being which had reasons for not preventing E, we would not be able to understand those reasons, then it seems reasonable to suppose that we shall not arrive at the conclusion that it is *very* likely that we shall fail to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E.¹²

In our view—and, we believe, in the view of most who now write on this topic¹³—our failure to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E *does* weakly support the claim that there are no such reasons. Thus, in our view—though perhaps not also in the view of most who now write on this topic—our failure to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E *does* weakly support the claim that there is no perfect being.¹⁴

¹² Here is another way of thinking about these issues. There are two ways in which it could be that there are reasons why a perfect being would not non-arbitrarily improve the world if it prevented E. First, there might be reasons which we are able to comprehend; second, there might be reasons which we are unable to comprehend. If we look for reasons, and are unable to find them, then our confidence in the first of these alternatives should be reduced; in consequence—unless our failure to find reasons should make us more confident that there are reasons which we are unable to comprehend—our confidence that there are reasons *tout court* should also be reduced. This reasoning can be blocked if we suppose that there *could not* be reasons that we are able to comprehend; but that is surely an incredible assumption to make.

It should perhaps be noted that we can't suggest that the reasons which a perfect being would have for not preventing E need not turn on the issue of whether the world would be non-arbitrarily improved by its prevention of E. For, while this suggestion—if cogent—would provide a reason for rejecting Wykstra's contention, there are good reasons for refusing to accept the suggestion. If 'the world'—the entire sphere of contingency—is not made better by the perfect being's prevention of E, then the perfect being simply does *not* have a reason to prevent E. This point—about the logical connection between reasons and the good—cannot be undermined by considerations about the vast gulf between the understanding of a perfect being and our own imperfect understanding.

¹³ Several people have suggested to us that Plantinga is one of those who disagrees, and that we really ought to say something about his views in this context. We agree that we ought to say something about Plantinga's views—though we prefer to defer that discussion to some other occasion—but we are not convinced that Plantinga does disagree. At any rate, a careful reading of [Plantinga 1988] did not make it obvious to us that Plantinga outright dismisses the suggestion that, other considerations aside, our failure to find reasons why the world would not be non-arbitrarily improved if a perfect being prevented E does make it less likely—however minutely!—that there is a perfect being. Plantinga [1988] makes it quite clear that he thinks that he could accept this claim with equanimity; moreover, [Plantinga 1988] also makes it quite clear that he would reject the corresponding claim about substantial support. But we have not been able to make an accurate determination of his view on the question which is the focus of the present section of our paper.

¹⁴ Perhaps those who think otherwise are misled by the thought that we should hardly expect to have insight into the reasons—the reasoning and motivation—of a perfect being. But, even if it is true

III

Even if it is granted that our failure to find reasons why a perfect being would not non-arbitrarily improve the world if it prevented E does provide weak support for the claim that there is no perfect being, it remains to be determined how strong this support is. It seems quite implausible to suppose that, all by itself, this failure is sufficient to establish that it is more likely than not that there is no perfect being. After all, there might be other reasons—perhaps even quite strong reasons—for believing that there is a perfect being. In order to assess this question, we need to take all of the relevant evidence into account—and that is not a task for a short discussion piece!¹⁵

However, even if we can't hope to show that our failure to find reasons why a perfect being would not non-arbitrarily improve the world if it prevented E is sufficient to establish that it is more likely than not that there is no perfect being, we might nonetheless hope to show that our failure to find reasons why a perfect being would not non-arbitrarily improve the world if it prevented E does provide *significant* support for the claim that there is no perfect being. While we concede that we might do more to explain what is required in order for one proposition to provide significant support for another, we shall follow the standard practice of hoping that this notion is sufficiently well-understood to allow us to proceed.

Given the way in which we have set up the discussion, the central task is to determine whether it is likely that, were there reasons why a perfect being would not non-arbitrarily improve the world if it prevented E, we would be able to find those reasons. If we suppose that it is quite likely—or very likely—that, *were there reasons why a perfect being would not non-arbitrarily improve the world if it prevented E, we would be able to find those reasons*, then we should hold that our failure to find reasons why a perfect being would not non-arbitrarily improve the world if it prevented E does provide significant support for the claim that there is no perfect being. On the other hand, if we deny that it is fairly likely—or very likely—that, *were there reasons why a perfect being would not non-arbitrarily improve the world if it prevented E, we would be able to find those reasons*, then we should hold that our failure to find reasons why a perfect being would not non-arbitrarily improve the world if it prevented E does not provide significant support for the claim that there is no perfect being.¹⁶ Why might we *deny* that it is quite likely that, were there reasons why a perfect

¹⁴ *continued*

that we should think it likely that our insight into the reasoning and motivation of a perfect being is limited, that is not relevant to the assessment of our argument. For the underlying question is whether the world would not be non-arbitrarily improved if a perfect being prevented E; and what we are looking for is reasons for thinking that this question should be answered in the affirmative. Thus—as much of the more recent literature recognizes—the central questions really concern limitations on our abilities to recognize goods rather than limitations on our insights into the reasoning and motivation of perfect beings.

¹⁵ Of course, there are questions about the prior probability that is assigned to R that might be raised at this point. If $\text{Pr}_{\text{prior}}(\text{R})$ is very much less than $\text{Pr}_{\text{prior}}(\sim\text{R})$, then more than weak support from F would be required to make $\text{Pr}_{\text{posterior}}(\text{R}) > \text{Pr}_{\text{posterior}}(\sim\text{R})$. But this is not the only way in which 'strong support' might be understood. If, for example, $\text{Pr}_{\text{prior}}(\text{R})/\text{Pr}_{\text{posterior}}(\text{R})$ is very much less than one, then there is a sense in which F gives strong support to R, even if it is also the case that $\text{Pr}_{\text{posterior}}(\text{R})$ is very much less than $\text{Pr}_{\text{posterior}}(\sim\text{R})$. We shall not need to worry about these kinds of niceties in what follows.

¹⁶ Note that the denial in question can take one of two forms. One could insist that it is fairly unlikely—or very unlikely—that, were there reasons why a perfect being would not non-arbitrarily

being would not non-arbitrarily improve the world if it prevented E, we would be able to find those reasons?

Well, as Bergmann [2001] argues, this judgement might follow from a more general scepticism about our knowledge of the realm of value. If we pay due attention to our cognitive limitations, and to the vastness and complexity of reality, then it may well seem plausible to suppose that our understanding of the realm of value falls miserably short of capturing all that is true about that realm. More positively, that same due attention might suggest that we are entitled to accept the following kinds of claims:

(ST1) We have no good reason for thinking that the possible goods we know of are representative of the possible goods there are.

(ST2) We have no good reason for thinking that the possible evils we know of are representative of the possible evils there are.

(ST3) We have no good reason for thinking that the entailment relations we know of between possible goods and the permission of possible evils are representative of the entailment relations there are between possible goods and the permission of possible evils.¹⁷

But, if we accept these kinds of claims, then won't we be entitled to claim that we are in no position to judge that it is fairly likely—or very likely—that, were there reasons why a perfect being would not non-arbitrarily improve the world if it prevented E, we would be able to find those reasons?

We think not. Suppose we take seriously the idea that it follows from our acceptance of ST1–ST3 that is not unlikely that there are goods beyond our ken—or relations beyond our ken between goods and evils (which themselves may or may not be beyond our ken)¹⁸—which justify a perfect being in not preventing E. Suppose further that we are, right now, witnesses to E, and that we could intervene to stop it at no personal cost. What we have just conceded is that, merely on the basis of our acceptance of ST1–ST3, we should insist that it is not unlikely that there is some good which, if we were smarter and better equipped, we could recognize as a reason for a perfect being's not intervening to stop E.¹⁹

¹⁶ *continued*

improve the world if it prevented E, we would be able to find those reasons. Alternatively, one could insist that we are in no position to form any judgement at all about the likelihood that, were there reasons why a perfect being would not non-arbitrarily improve the world if it prevented E, we would be able to find those reasons. While either of these positions would suffice for the purposes of sceptical theists, it is the latter which shall be the major focus of our attention.

¹⁷ There is potential ambiguity in ST1–ST3. We take it that the right way to read, say, (ST1) is as follows: we have no good reason for thinking that the goods we know of are representative of the goods that there are in the world. It is less controversial that we have no good reason for thinking that the goods we know of are representative of the goods that there are in all possible worlds. But that less controversial claim is of no use to the sceptical theist, since—as Rowe insists—goods in other possible worlds cannot justify God's actions in our world.

¹⁸ Hereafter, we shall use the shorthand formulation which refers only to goods beyond our ken, and omits reference to the considerations about relations beyond our ken between goods and evils. However, this is only for ease of formulation: it is very important—as Bergmann and Alston both insist—that it is not merely the goods and evils which may be beyond our ken.

¹⁹ Perhaps it might be disputed that we did previously concede this. If we suppose that we are essentially unable to recognize the reasons in question, then—no matter how much smarter and better equipped we were—we would not be able to recognize the reasons for the perfect being's not intervening to stop E. However, there is nothing other than ease of exposition which is lost if we grant this point—parity of reason *will* still get us to the claim that, for all we know, it would be best, all things considered, if we did not intervene.

Plainly, we should also concede—by parity of reason—that, merely on the basis of our acceptance of ST1–ST3, we should insist that it is not unlikely that there is some good which, if we were smarter and better equipped, we could recognize as a reason for our not intervening to stop the event. That is, our previous concession surely forces us to allow that, given our acceptance of ST1–ST3, it is not unlikely that it is for the best, all things considered, if we do not intervene. But, if we could easily intervene to stop the heinous crime, then it would be *appalling* for us to allow this consideration to stop us from intervening. Yet, if we take the thought seriously, how can we also maintain that we are morally required to intervene? After all, as a result of our acceptance of ST1–ST3, we are allegedly committed to the claim that it is not unlikely that it would be for the best, all things considered, if we did not do so.

Bergmann claims that we should give high probability to ST1–ST3. What seems to follow from this—at least if we follow the model which Bergmann offers in the case of the claim that there are goods which, if we were smarter and better equipped, we would recognize as reasons for *a perfect being* not to prevent E—is that we have no good reason to assign a low probability to the claim that there are goods which, if we were smarter and better equipped, we would recognize as reasons for *us* not to prevent E. Yet, if we do not have good reason to assign a low probability to the claim that there are goods which, if we were smarter and better equipped, we would recognize as reasons for us not to prevent E, then how can we have good reason to interfere and to prevent it? True enough, rape and murder are terrible evils; hence, that a particular event is an instance of rape and murder is a *pro tanto* reason to prevent that event. But an event can be many things at once, and even a rape and murder *could* be an event which is a very great good (e.g., the prevention of the destruction of the world, to take a hackneyed example). If we are not prepared to judge that it is unlikely that a particular instance of rape and murder is not also a very great good—and that is just the kind of judgement which acceptance of ST1–ST3 is supposed to preclude—then we do not have sufficient reason to interfere, and to prevent the rape and murder, no matter how little it would cost us to do so.

The conclusion for which we have argued here is that what Bergmann calls ‘sceptical theism’ really does involve an unacceptable scepticism if it is strong enough to provide a telling objection to evidential arguments from evil. Bergmann considers some other ways in which one might try to make this case and (rightly) argues that they fail.²⁰ However, our discussion seems to show *either* that our moral practice implicitly commits us to the assignment of a high probability to the claims which, according to ST1–ST3, we have no good reason for believing; or else that ST1–ST3 are insufficient to license the conclusion that it is not unlikely that there are goods beyond our ken which would justify a perfect being in not preventing E. Perhaps there is some reason to think that ST1–ST3 embody a relatively benign form of scepticism—perhaps it is plausible that the goods we know fail to be representative of the goods there are; that the evils we know fail to be representative of the evils there are; and that the entailment relations we know of between goods and the permission of evils are representative of the entailment relations there are between goods and the permission of evils—but, if so, then there must be some problem with the inference which Bergmann draws. Our moral practice—our ordinary moral behaviour—shows that we do think it unlikely that there are goods beyond our ken which would justify us in not

²⁰ We shall return to some of Bergmann’s arguments in Section V.

preventing E; so there is plainly room for serious doubt about the suggestion that considerations like ST1–ST3 are sufficient to establish that it is not unlikely that there are goods beyond our ken which would justify a perfect being in not preventing E.

Here is another way of making our key point. Suppose that we try to give a rational reconstruction of the moral reasoning that we undertake when we reach the decision to intervene in the case in which we can easily prevent rape and murder. The reconstruction will have to go something like this:

- (1) There is *pro tanto* reason for me to intervene to prevent E. (Indeed, I have a *pro tanto* duty to intervene to prevent E.) (Premise)
- (2) I have found no *pro tanto* reason for me not to intervene to prevent E. (Premise)
- (3) (Hence) There is no *pro tanto* reason for me not to intervene to prevent E. (From 2)
- (4) (Hence) I have all things considered reason to intervene to prevent E. (From 1, 3)

If we like, we can make this reconstruction look even more like the evidential argument from evil, by casting it in terms of reasons why the world would not be non-arbitrarily improved if I were to prevent E. However, even the version which we have given makes the point clearly enough: our reasoning from *pro tanto* reasons to all things considered reasons always relies upon a ‘noseeum’ inference of just the kind which appears in our evidential argument from evil. If sceptical theism is sufficient to block ‘noseeum’ inferences about values, then we lose our ability to reason to all things considered conclusions about what to do.

If the case which we have argued is cogent, then the sceptical theist opponent of the evidential argument from evil has been placed in an uncomfortable position. The sceptical theist wants to be able to claim that it is not unlikely that there are unknown goods which would justify a perfect being in not preventing E. Yet, if the considerations to which the sceptical theist appeals can establish this, then they will also suffice to establish that it is not unlikely that there are unknown goods which would justify us in not preventing E. But, if we do believe that it is not unlikely that there are unknown goods which would justify us in not preventing E, then it is very hard to see how we could fail to be justified in not preventing E.

The key question, on our view, is not—as sceptical theists have typically supposed—whether there is reason to hold that it is not unlikely that there are goods outside our ken which would justify a perfect being in not preventing E; rather, the key question is whether there is reason to hold that it is not unlikely that there are goods outside our ken which would justify a perfect being in not preventing E *which is not also* reason to hold that it is not unlikely that there are goods outside our ken which would justify us in not preventing E. Perhaps unsurprisingly, most extant discussions in the literature are little help on this point. For instance, Howard-Snyder [1996c] gives two arguments for the claim that it is fairly likely that there are goods outside our ken which would have a role to play in the deliberations of a perfect being. However, even if those arguments—the *Progress* argument and the *Complexity* argument—are good reasons for thinking that there are goods outside our ken which have a role to play in the deliberations of a perfect being, those arguments are plainly of no help at all in establishing the conclusion that, while it is very unlikely that there are *unknown goods* which would be forgone if we were to prevent E, it is fairly likely that there are *unknown goods* which would be forgone if a perfect being were to prevent E. Yes—perhaps!—facts about axiological discoveries and the complexity of moral considerations make it likely that there are goods outside our ken which would not be outside a perfect being’s ken; but those facts do not lend *any* significant support to the claim that,

while it is very unlikely that there are unknown goods which would be forgone if we were to prevent E, it is not very unlikely that there are unknown goods which would be forgone if a perfect being were to prevent E.

Of course—as we noted at the beginning of our discussion—it might be that perfect being theists have independent reasons for thinking that it is fairly likely that there are unknown goods which would be forgone if a perfect being were to prevent E; so we do not suppose—even if the considerations which we have urged thus far are cogent—that we have shown that the evidential argument from evil does succeed in establishing, or even in strongly supporting, the claim that there is no perfect being. However—unless we have made a mistake elsewhere in our assessment of the argument—it seems to us that we have now done enough to show that Bergmann’s version of ‘sceptical theism’ does not provide a good response to the claim that the evidential argument from evil provides significant support for the claim that there is no perfect being.

IV

Where might we have gone wrong in our argument? Well, one obvious suggestion is that we have not paid due attention to the fact that it makes a huge difference that it is not we but rather the perfect being who fails to prevent E. But what reason is there to believe that this difference is so important? If there are goods utterly beyond our ken—and if the motives and purposes of a perfect being are utterly beyond our ken—then it is hard to see that we have any reason to think that this is a relevant difference. Given that we have no knowledge of the goods at issue, surely it must be an entirely open question whether they can be secured by our failure to prevent E.

Perhaps it might be replied: we *know* that a perfect being would not make a world like *that*; indeed, we know that a perfect being would set things up so that great goods will be secured by our prevention of E. Really? And how is this alleged knowledge supposed to be compatible with our unutterable ignorance about the motives of a perfect being, and the goods that there might be? How do we get to know *these* things, while remaining ignorant about those other things that we are required not to know in order to have a reply to the evidential argument from evil? Surely, given what sceptical theists claim can be made of considerations like ST1–ST3, they are obliged to maintain that we can assign no probability to the judgement that a perfect being would not make a world in which very great goods are secured by our failure to prevent E.

Perhaps it might be replied: there is a perfect being, and that being has revealed to us—by way of personal experience, scripture, and religious tradition—that the world is not set up in that way. By this kind of communication, the perfect being guarantees—or, anyway, makes it overwhelmingly likely—that there are no goods to be secured by our failure to prevent E. (Actually, the alleged commandments are typically stronger than that: we have a guarantee that the perfect being wants us to prevent events like E when they threaten to occur.) *But*, of course, in making this reply, one would be giving up the sceptical theist ambition: it is no longer true that it is merely considerations about our cognitive limitations that yield the desired conclusion. The whole point of the sceptical theist response is that it is supposed to avoid appeal to the other evidence which theists possess for the existence of a perfect being and the directives that that being makes in connection with our behaviour.

Perhaps it might be replied that there are significant relevant differences between what would be good for our creator to do and what would be good for us to do. Parents have certain rights over their children which other people do not. This arises from their being (to a limited extent) the source of their children's existence and well-being. Hence, plausibly, parents have the right to send children to a neighbourhood school at which they will not be totally happy, for the benefit of the community as a whole. Strangers do not have those sorts of rights over our children. Plausibly, the state has certain rights over us that we do not have over each other—again, in virtue of its being a benefactor. If this is right, then a perfect being has far greater rights—in view of being a far greater benefactor—to allow us to suffer for the common good (e.g., to endure rape and murder), whereas we do not have those sorts of rights.

Many variants of this objection can be imagined. All seem to us to suffer from the same difficulty, namely, that the fact that there are differences between us and a perfect being with respect to goods of which we have knowledge is completely beside the point. The sceptical theist claim is that, in light of considerations like ST1–ST3, we have good grounds for holding it not unlikely that there are *unknown* goods which are secured by the failure of a perfect being to prevent E. Our claim is that, if this is right, then surely those same considerations give us good grounds for holding it not unlikely that there are unknown goods which will be secured by our failure to prevent E. Of course, we do not claim that it is the same unknown goods in each case (nor, even, do we claim that this is likely): that there are different agents involved may well make this further claim implausible. However, if we accept that—on the grounds given by Bergmann—it is not unlikely that *unknown goods* are secured by a perfect being's failure to prevent E, how can we deny that—on the grounds given by Bergmann—it is not unlikely that *unknown goods* will be secured by our failure to prevent E?

For the reason just given, we think that no appeal to the enormous differences between us and a perfect being can have any effect on our argument. So long as sceptical theists appeal to considerations which rule out all 'noseeum' inferences about values if they rule out any, their position will surely be vulnerable to the kind of reply which we have given.

V

What else might have gone wrong with our argument? Well, another fairly obvious suggestion is that, while, in fact, ST1–ST3 do justify us in accepting the conclusion that, for all we know, there are unknown goods which would be secured by our failure to prevent E, it does not follow that this is any threat to our ordinary moral practice.

We think that this suggestion is quite implausible: if ST1–ST3 do justify us in accepting the sceptical theist conclusion then—for the reasons given earlier—there is a massive impediment to our reasoning to the conclusion that we ought to try to prevent E. Of course, it is crucial not to mislocate the difficulty which we diagnose: we are not supposing that the sceptical theist is unable to assert that rape and murder are wrong, nor that the sceptical theist is unable to assert that we have a *pro tanto* duty to prevent rape and murder. Rather—to put the point a little tendentially—our claim is that if you refuse to make any positive judgements about likelihoods, then you are unable to reason your way to decisions. To explain this contention, we shall begin with a case that has nothing to do with values.

Suppose I am trying to decide whether or not to have Weetbix for breakfast. It is *possible* that something very momentous—say, the outcome of the next Federal election—turns on my choice. It is possible, too, that the relevant connections between my choice and the election outcome would be utterly obvious to a smarter and better informed creature while being utterly beyond my ken. However, when I engage in my deliberation, I—rightly—assign negligible probability to this possibility, and, in this way, am able to ignore it. Suppose, however, that I am sceptical about my entitlement to make probability judgements of this kind; suppose, that is, that I am not prepared to assign negligible probability—nor, indeed, any other kind of probability—to the claim that my having Weetbix for breakfast will determine that the Liberals win the next election. Then, it seems to us, I am not able to make a decision about what to do to which I can give first personal endorsement; whatever I do, I take seriously the idea that it might be the choice which determines the outcome of the next election to be the result which I desperately do not want. Of course, I can toss a coin—and, in that way, make a decision—but I cannot use the resources that I have at hand in order to reason my way to a choice. (Pretend that we are expected utility maximizers. I can only use the utility calculus to arrive at a decision if I am prepared to assign probabilities—or *proper* probability intervals—to outcomes; if some outcomes are simply assigned the interval $[0, 1]$, then—except in a small number of very special cases—calculation is stymied. Even for those who reject decision theory, this result ought to be suggestive.)

Now, of course, exactly the same point applies to values (and utilities). Suppose I'm trying to decide whether or not to perform a certain actions—say, (try to) prevent the occurrence of E. It is possible that there are goods and evils about which I know nothing which attach to my performance or non-performance of the action. (Note that we do not assume that the value of the action is determined by its consequences; the values in question can be of any kind whatsoever.) It is possible, too, that these goods and evils about which I know nothing would be utterly obvious to a smarter and better-informed creature. However, when I engage in my deliberation, I—rightly—assign negligible probability to this possibility and, in this way, am able to ignore it. Suppose, however, that I am sceptical about my entitlement to make probability judgements of this kind; suppose, that is, that I am not prepared to assign negligible probability—nor, indeed, any other kind of probability—to the claim that the total value of my permitting E is very high. Then, it seems to us, I am not able to make a decision about what to do to which I can give first personal endorsement; whatever I do, I take seriously the idea that the alternative course of action is the one which has the greater value. Of course, I can toss a coin—and, in that way, make a decision—but I cannot use the resources that I have at hand in order to reason my way to a choice. (Pretend, again, that we are expected utility maximizers. I can only use the utility calculus to arrive at a decision if I am prepared to assign values—or *proper* value intervals—to outcomes; if some outcomes are simply assigned the value interval $[-\infty, \infty]$, then—except in a small number of very special cases—calculation is stymied. Even for those who reject decision theory, this result ought to be suggestive.)²¹

²¹ Given the parenthetical asides at the conclusion of the present paragraph and the preceding paragraph, some readers may be tempted to suppose that our overall argument relies on the unjustified meta-ethical assumption that some version of consequentialism is correct. This supposition would be a mistake. As we noted earlier, our argument does not require us to suppose that the values of courses of action open to us are decided *merely* by the values of

In each case, there are alternatives to tossing a coin which amount to the same thing. I could, for example, ‘calculate’ by the goods of which I have knowledge, and in that way arrive at a ‘decision’. (And, in the former case, I could calculate according to the causal considerations of which I have knowledge.) But, from the standpoint of everything that I believe and have reason to believe, these really are arbitrary choices; for, by hypothesis, I am supposed to be completely agnostic about the probability and value assignments which are assumed in order to make a decision possible.

If the forgoing argument is right, then there is a clear sense in which the sceptical theist argument against the evidential argument from evil, if cogent, does undermine our ordinary moral practice; for, if cogent, it undermines our ability to engage in perfectly ordinary kinds of moral reasoning. However, before we close this section, we shall consider some further claims that have been made on behalf of sceptical theism in connection with the claim that it undermines ordinary moral practice.

Bergmann argues for the claim that the sceptical theist argument poses no threat to ordinary moral practice on the grounds that one can quite consistently refuse to hold that it is unlikely that there are unknown goods secured by our refusal to prevent E, while nonetheless maintaining that our refusal to prevent E is morally wrong. For example, it could be that while goods and evils about which one knows give rise to *pro tanto* duties, unknown goods and evils do not. Hence, a perfect being, which knows all goods and evils, has an all things considered duty to permit E, while we have an all things considered duty to prevent it.

The particular suggestion plainly won’t wash: even if we grant that unknown goods and evils cannot give rise to *pro tanto* duties, it is simply a mistake to suppose that the only factors which must be weighed when we determine what it is that we ought to do are our *pro tanto* duties. After all, apart from our *pro tanto* duties, we also have ‘*pro tanto* permissions’—rights, etc.—which need to be weighed in our deliberations. Moreover—and this is the crucial point—there is always a ‘noseeum inference’ which is required in order to reach a result from our deliberations: we have to be able to infer from the fact that we can see no other relevant considerations (about duties, rights, interests, etc.) that there are no other considerations which are relevant to our deliberations. If we cannot assign a low probability to the claim that we have missed some relevant considerations, then we cannot claim that our deliberations have resulted in a clear verdict. But, of course, unknown goods—and unknown implicational relations involving goods—can give rise to ‘permissions’ (rights, etc.): while, perhaps, I cannot have duties which I do not know about, I plainly can have rights about which I am completely in the dark. So, to return to the case at hand, I do not require a competing *pro tanto* duty not to intervene in order to be justified in permitting E; rather, all I require is that I fail to attribute a sufficiently low probability to the claim that I have ‘permissions’—rights, etc.—which outweigh my *pro tanto* duty to prevent E. If, for example, I cannot assign any likelihood to the claim that both I and the

²¹ *continued*

the consequences of those actions. We think it plausible that consequences are at least sometimes morally relevant considerations; but nothing in our argument requires us to assume even this much. No matter what view of values one takes, it is possible for one to be faced with a choice between a range of possible actions in circumstances in which one can do no more than attach probabilities to a range of possible future states of the world. In order to make decisions about what to do, one must then have some way of trading off competing values of actions—conflicting *pro tanto* duties, *pro tanto* rights, and the like—whether or not one thinks that the range of possible future states of the world has any bearing on the ultimate decision which one reaches.

victim will benefit enormously if I do not attempt to prevent E, then how can I assign any likelihood to the claim that I have a *pro tanto* ‘permission’ not to attempt to prevent E?

The reply which we have given to the particular suggestion carries over to Bergmann’s more general argument as well. If there is a person who refrains from intervening in the case of E because they have internalised the sceptical theist’s scepticism, and hence are unable to find a justification for intervening to which they can give first personal endorsement, then sceptical theists are simply in no position to say this person’s failure to intervene is wrong. Bergmann makes it easy for himself by following Russell [1996] in considering a case in which a person fails to prevent E and merely lacks an appropriate motivating reason for doing so. However, even ignoring the point that, at most, these considerations establish that there is a *pro tanto* case for claiming that what the person does is wrong—after all, by the lights of sceptical theists, we can assign no likelihood to the claim that there are unknown goods which justify what the person does!—the crucial observation to make is that the success of the sceptical theist objection to evidential arguments from evil gives such a person all the motivating reason that they need. Or—more exactly—if the person whom we are considering accepts the sceptical theist objection to evidential arguments from evil, and if we do so as well, then we have no basis for claiming that their inaction is wrong: by their lights and ours, their moral reasoning to the conclusion that they lack sufficient reason to prevent E is impeccable.²²

In sum, then, the sceptical theist response to evidential arguments from evil, if successful, really would pose a serious threat to ordinary moral practice. In any decision situation, we would be in the position of the person who is ‘out of her depth’ and who knows that she is ‘out of her depth’. What is objectionable about this is not the thought that we might always be ‘out of our depth’, in the sense that we are unable to *fully* evaluate the considerations which bear on our decisions: for, of course, none of us can know all of the long term consequences of any action we perform—no doubt there were all kinds of good deeds which were causally necessary for Hitler to be born—and to this extent we are always ‘out of our depth’ in deciding what to do. Rather, the problem is that, if we are always ‘out of our depth’ *and* if we are always aware that we are ‘out of our depth’, then we can never give first personal endorsement to any of our actions; moral deliberation can never end in anything more than the equivalent of tossing a coin. Unless we can give *sufficient* evaluation of the considerations that bear upon our decisions, we are not capable of reaching any embraceable moral decisions at all.

VI

Suppose that our argument against Bergmann [2001] is good. Plainly enough, it doesn’t follow immediately that no ‘sceptical theist’ response to evidential arguments from evil is good. After all, for all that we have argued so far, it could be that Bergmann rather incautiously defends a more extreme position than ‘sceptical theists’ are normally inclined

²² Part of our complaint against Bergmann is that he simply assumes that one must have a positive motivating justification in order to perform an action, e.g., permitting E. But one can be justified in performing an action if the outcome of one’s deliberation ends in deadlock; Buridan’s ass is justified in eating the left bale of hay as a result of an arbitrary choice. Similarly, one influenced by the sceptical theist argument could be justified in permitting E as the result of an arbitrary choice, despite the absence of any more positive justification for the action.

to defend. Perhaps, when we look at the writings of Alston, Howard-Snyder, van Inwagen, Plantinga, Wykstra, and others, we shall find that the views which they defend are more subtle and nuanced than the extreme position that is defended by Bergmann.

And perhaps not. Since we don't have space to examine all of the remaining 'sceptical theists' here, we shall try to make our case only in connection with Alston [1991, 1996]; however, we think that the generalization from what we have to say about Alston is pretty straightforward. What we propose to argue is that, in fact, the position that Alston defends is open to exactly the same kind of objection that we have levelled at Bergmann.

Alston [1996] claims that it is evidently absurd to suppose that the fact that we cannot see what sufficient justifying reason an omniscient, omnipotent being may have for doing something provides strong support for the supposition that no such reason is available for that being. In support of this claim, he mentions two major kinds of considerations. First, it just seems wrong to take the insights attainable by finite fallible human beings as an adequate indication of what is available in the way of reasons to an omniscient, omnipotent being. Second, when we look for justifying reasons in the kind of case to which Rowe adverts, we are involved in an attempt to determine whether there is a so-and-so in a territory the extent and composition of which is largely unknown to us. (Alston [1991] provides the following further claims about this territory: (i) we lack data; (ii) we face greater complexity than we can handle; (iii) we face difficulties in determining what is possible; (iv) we are ignorant of the full range of possibilities; (v) we are ignorant of the full range of values; and (vi) there are limits to our capacity to make well-considered value judgements. Alston [1996] emphasises that it is not merely that there may be goods and evils which are unknown to us: it may also be that our grasp of the nature of goods known to us is not sufficient for us to properly assess their degree and kinds of value; and it may be that our grasp of the conditions of the realization of known goods is insufficient for us to properly assess what would be required in order for a perfect being to bring about those goods.)

Alston's view, then, is that when we consider human cognitive limitations in a variety of spheres including the realm of value, we see that we have no reason to attribute substantial credence to the claim that, if there were reasons why a perfect being would not non-arbitrarily improve the world if it prevented E, we would be able to find those reasons. Hence, in particular, we have no reason to attribute substantial credence to the claim that, if there were goods—or connections between goods and evils—capable of justifying a perfect being in permitting E, then we would be able to find those goods—or connections between goods and evils. Moreover, Alston is also committed to the claim that, in consequence, this consideration of human cognitive limitations in a variety of human spheres including the realm of value is alone sufficient to block the inference from the claim that we can find no goods—or connections between goods and evils—that would justify a perfect being in permitting E to the conclusion that there are no goods—or connections between goods and evils—that would justify a perfect being in permitting E.

But, if that is right, then it is clear that the argument that we deployed against Bergmann can also be deployed here. If those considerations about human cognitive limitations are alone sufficient to block the inference from the claim that we can find no goods—or connections between goods and evils—that would justify a perfect being in permitting E to the conclusion that there are no goods—or connections between goods and evils—that would justify a perfect being in permitting E, then those considerations are also alone

sufficient to block the inference from the claim that we can find no goods—or connections between goods and evils—that would justify us in permitting E to the conclusion that there are no goods—or connections between goods and evils—that would justify us in permitting E. (As we noted above, Alston relies on the claim that we are ignorant of the full range of goods and evils; that we are ignorant of the full range of connections between goods and evils; that we lack data; that we face greater complexity than we can handle; that we are ignorant of the full range of possibilities; and that there are limits to our capacity to make well-considered value judgements. But if *these* claims are enough to block the inference in Rowe's argument, then surely they must be enough to block the corresponding inference in the case in which we are deciding whether to intervene to prevent E.)

Perhaps the position that Alston defends can also be attacked from another direction. Suppose that we agree that we can assign no probability to the claim that there is some great good that would be lost if a perfect being were to intervene to prevent the occurrence of E. How then can we pretend to be able to assign any probability to the claim that there is some great good that would be lost if *we* were to intervene to prevent the occurrence of E? Perhaps—for all we know—there are very great agent-centred goods that will be lost if we intervene. (Perhaps, for example, the agent will be denied salvation as a result of our intervention. Given our alleged ignorance about the motives of a perfect being, who can say?) Perhaps—for all we know—there are very great non-agent-centred goods that will be lost if we intervene. (Perhaps, for example, there are outweighing benefits that accrue to the attacker and to bystanders in the next life that will be lost if we intervene.²³ Again, given our alleged ignorance about the motives of a perfect being, who can say?) If we really—seriously!—think that we can assign no probability to the claim that our intervention to prevent E would disbar the victim from entry to heaven, then surely we have no way at all of making a reasoned choice about whether or not to intervene. Yet, if we are so ignorant about the motives of perfect beings, then how can we assign any probability to the claim in question?

VII

In conclusion, it may be useful to discuss some of the ways in which our argument may be misunderstood. The comments of one of our referees can serve as a useful foil here. The referee writes:

The authors argue against Bergmann to the effect that if his sceptical position has the intended force against the argument from evil, it will have a debilitating force on ethical reasoning generally. This argument seems to me to neglect a crucial difference between speculative and practical reasoning. The authors claim that if we cannot know that God does not have a good reason to prevent a certain horrific evil, then we can't know that there is not a good reason for us to avoid preventing it even if we could. But the former point concerns an argument against the existence of God, a speculative (theoretical) activity in which we have no practical concern. We can't do anything

²³ Of course, were this so, it would be 'just one more instance of evil the reason for permitting which we cannot see'. (Cf. Alston [1996: 321].)

about whether God exists, whether he prevents something, or what reasons he does or doesn't have. But the other point involves reflection on whether we should perform a certain action that is open to us. And here the lack of knowledge has a different bearing. Where we are faced with choice to perform A or not, we are, if we make a rational decision, forced to act on the best relevant knowledge we have. Even if we don't know that there isn't a good reason to abstain, it would be irrational to abstain from preventing the evil because of that epistemic possibility. Here we are forced, if we are rational, to make our decision on the basis of what we know (and rationally believe). Hence it would be foolish to avoid action because of something we don't know not to obtain. Whereas in the question of the existence or goodness of God, we are not called upon to make a practical decision to perform an (overt) action but only to evaluate an argument.

First, what we don't say. We *don't* argue that, if we can't know that there is not a good reason for us to avoid preventing E, then we have a reason not to prevent E. We *don't* argue that, if it is epistemically possible that there are outweighing goods that will be lost if we prevent E, then we have a reason not to prevent E. There are at least two difficulties here. First, the key questions are not particularly questions about knowledge (nor even, perhaps, rational belief). When faced with the decision about whether or not to perform an action, rational people act on the basis of the relevant considerations that are available to them (whether or not those considerations count as knowledge or even, perhaps, rational belief). Second, and more importantly, we have to be able to factor self-confessed ignorance into our process of deliberation. It need not be 'foolish to avoid action because of something we don't know not to obtain'. Suppose that we don't know that it is not the case that the floorboards are rotten (or, taking account of the previous point, that we are not prepared to make any estimation of whether or not it is the case that the floorboards are rotten). How can we then step with complete confidence onto the floor? In practical considerations—no less than in theoretical considerations—probabilities are the stuff of deliberation. And, whereof one is not prepared to assign probabilities, thereof one is simply not able to deliberate. (Perhaps it is also worth noting that we object to the suggestion that 'lack of knowledge' has a different bearing in the case of theoretical reason than it does in the case of practical reason. True enough, in the case of theoretical reason one can suspend belief; whereas, in the case of practical reason, one must end with a 'decision' about what to do. However, the key point is that 'deliberation' can run out in just the same kind of way in both cases: because of the imperative to act, one may perform the equivalent of tossing a coin—something that would be reprehensible in the case of theoretical inquiry—but it should not be supposed that this counts as reasoning one's way to a decision.)

Second, what we do say. Our central claim is that, if the considerations to which 'sceptical theists' appeal—considerations of human cognitive limitations in the realm of values (and perhaps elsewhere as well)—were alone sufficient to undermine the noseem inference in evidential arguments from evil, then those considerations would also be alone sufficient to undermine familiar and ordinary kinds of moral reasoning. If the kinds of considerations to which sceptical theists appeal entail that we can assign no probability to the claim that there are great goods that are secured by the failure of a perfect being to prevent E, then the kinds of considerations to which sceptical theists appeal also entail that we can assign no probability to the claim that there are great goods that are secured by our

failure to prevent E. But if we can assign no probability to the claim that there are great goods that are secured by our failure to prevent E, then we cannot arrive at a reasoned view about whether or not to intervene to prevent E. And that's not an acceptable result.

University of Texas at San Antonio
Monash University

Received: August 2002

Revised: March 2003

REFERENCES

- Alston, William 1991. The Inductive Argument From Evil and the Human Cognitive Condition, in *Philosophical Perspective 5: Philosophy of Religion*, ed. James Tomberlin, Atascadero CA: Ridgeview: 29–67.
- Alston, William 1996. Some (Temporarily) Final Thoughts on Evidential Arguments from Evil, in Howard-Snyder 1996b: 311–32.
- Bergmann, Michael 2001. Sceptical Theism and Rowe's New Evidential Argument From Evil, *Noûs* 35: 278–96.
- Fitzpatrick, F. J. 1981. The Onus of Proof in Arguments about the Problem of Evil, *Religious Studies* 17: 19–38.
- Howard-Snyder, Daniel 1996a. Introduction, in Howard-Snyder 1996b: xi–xx.
- Howard-Snyder, Daniel, ed. 1996b. *The Evidential Argument from Evil*, Bloomington: Indiana University Press.
- Howard-Snyder, Daniel 1996c. The Argument from Inscrutable Evil, in Howard-Snyder 1996b: 286–310.
- van Inwagen, Peter 1991. The Problem of Evil, the Problem of Air, and the Problem of Silence, in *Philosophical Perspective 5: Philosophy of Religion*, ed. James Tomberlin, Atascadero CA: Ridgeview: 135–65.
- Plantinga, Alvin 1979. The Probabilistic Argument from Evil, *Philosophical Studies* 35: 1–53.
- Plantinga, Alvin 1988. Epistemic Probability and Evil, in Howard-Snyder [1996b]: 69–96.
- Rowe, William 1979. The Problem of Evil and Some Varieties of Atheism, *American Philosophical Quarterly* 16: 335–41.
- Russell, Bruce 1989. The Persistent Problem of Evil, *Faith and Philosophy* 6: 121–39.
- Russell, Bruce 1996. Defenceless, in Howard-Snyder [1996b]: 193–205.
- Wykstra, Stephen 1984. The Humean Obstacle to Evidential Arguments from Suffering: On Avoiding the Evils of 'Appearance', *International Journal for Philosophy of Religion* 16: 73–93.
- Wykstra, Stephen 1996. Rowe's Noseeum Arguments From Evil, in Howard-Snyder [1996b]: 126–50.