

6

The Father of Lies?

Hud Hudson

1. The Elmer Fudd Problem

My garden isn't doing very well. I suspect there is a rabbit in the garden. I'd like to know. I call the A₁-Rabbit Identification and Removal Service. They come out, and after 20 minutes of sitting on the edge of the garden, having some lunch, and stamping a bit here and there, they knock at the door ready to be paid:

"No rabbits here." "Really? How do you know?" "Look, if there were a rabbit in the garden, then we would be aware of it and we would recognize it as a rabbit, but we aren't aware of any rabbits in the garden." "Well... why should I accept the conditional that takes you from admitted ignorance to confident verdict?" "Because even though we haven't searched every possible rabbit-hiding place in the garden, we have searched a couple of them and we came up empty."

I don't pay. I like the move, though. The A₁-Team took a sample of the garden, found it empty, engaged in inductive generalization, and adopted their no-rabbit stance. But their sample was unrepresentative—in this case because it was abysmally small—and that fact adequately undermines the strength of the inductive inference they have offered me.

I'd still like to know the answer to my question. I call the A₂-Rabbit Identification and Removal Service. They come out, and after 20 hours of exhausting rabbit-searching, they knock at the door ready to be paid:

"No rabbits here." "Really? How do you know?" "Look, if there were a rabbit in the garden, then we would be aware of it and we would recognize it as a rabbit, but we aren't aware of any rabbits in the garden." "Well... why should I accept the conditional that takes you from admitted ignorance to confident verdict?" "Because we

have searched every possible rabbit-hiding place in the garden and have done so in such a way that a rabbit cannot hide in any once-searched spot, and we came up empty.”

I pay. Unlike before, I’m not offered a defense of the crucial conditional by way of inductive generalization, but rather by way of exhaustive and complete examination—no turnip unmolested, no lettuce unturned. (And—*take note*—I simply spotted the A2-Team thesis that they know a rabbit when they see it.)

The relevance of our two warm-up cases:

The theist and atheist can often come to agreement about a claim of this sort: “Event E is evil, and given God’s essential omnipotence, essential omniscience, and essential perfect goodness, God would have prevented the occurrence of E unless there were a compensating good or some other morally justifying reason for the permission of E.” Theist and atheist then part company on the question of whether there is such a compensating good or morally justifying reason.

Moreover, the theist and atheist can often come to agreement about a further claim, as well: “We are not aware of any compensating good or morally justifying reason for this or that (admittedly) evil state of affairs.” A popular atheistic move at this juncture is to introduce a bridge premise that exploits some such inscrutable evil, E, a bridge premise of the following form: “If there were a compensating good or some other morally justifying reason for E, then we would be aware of it and able to recognize it as such.” Once this bridge premise is defended, the joint admission of ignorance can be used to infer that there is no compensating good or other morally justifying reason for E. Accordingly—given the first point of agreement noted above—the inference to atheism is secured.

It is worth noting that with respect to compensating goods or morally justifying reasons we can claim (at the very least) to have matched the efforts of the A1-Team. That is, we have stamped around a bit in value theory and in a theory of permissions, and we have disqualified a few candidates for being the morally justifying reason for the permission of world’s horrific evils. But if our sample should turn out to be unrepresentative (perhaps by being abysmally small)—we shouldn’t pay. We’ll need a strong inductive argument for that bridge premise, and that depends on just how much more we’ve accomplished in our search than the A1-Team accomplished in theirs.

It is also worth noting that we certainly cannot claim to have matched the efforts of the A2-Team. The garden of abstracta (where hide the potential compensating goods and other morally justifying reasons) is infinitely large. Whereas we might be inclined to pay if we had surveyed the lot—we quite simply haven't; in fact we've barely begun.

A far more interesting case:

The following year (when the garden is ailing once again) I'd like myself another rabbit report... but I don't want to pay the exorbitant prices of the A2-Team. I call the A3-Rabbit Identification and Removal Service. They come out, and after 12 hours of rigorous and systematic searching, they knock at the door ready to be paid:

"No rabbits here." "Really? How do you know?" "Look, if there were a rabbit in the garden, then we would be aware of it and we would recognize it as a rabbit, but we aren't aware of any rabbits in the garden." "Well... why should I accept the conditional that takes you from admitted ignorance to confident verdict?" "Because even though we haven't searched every possible rabbit-hiding place in the garden, we have searched a very large and representative portion of it—some 60% in fact, and we came up empty. That's good enough for us."

It's good enough for me, too, and I start to write that check. As did the hopeless A1-Team, the A3-Team has offered me an inductive generalization, but unlike their undistinguished predecessors, they have produced an inductive generalization allegedly supported by a representative sample of the whole. But then, in the middle of writing that check, I find myself worrying about a couple of issues.

First, a small point: Whether or not the sample is representative will not simply be a function of its size, even if it is very large. And, of course, the analog I care about—the garden of abstracta—is so infinitely vast that sheer sample size can't be reported with phrases like "we have surveyed 60% of the candidates for morally justifying reasons."

Second, a somewhat larger point: We are interested in whether we have good reason to believe the bridge premise that promises to take us from an admission that we do not know of a morally justifying reason to the conclusion that there isn't one. Suppose we wanted to undermine the strength of the inductive generalization described above which purports to furnish us with that good reason.

Accordingly, consider a continuation of the conversation at the door in our A3-Rabbit Identification and Removal Service Case:

“I know you boys searched a very large portion of the garden, but you’ll note that a significant portion of the garden is covered by a tarp held in place by rocks and earth. Was any of your sample taken from under the tarp?” “No, sir, we gave it a try, but that tarp was just too heavy to lift.” “But given the business you’re in, I’m sure you’ll agree that it’s likely that if there were a rabbit in the garden, it would be hiding under the tarp.” “Yes, sir.” “Then, despite its being quite large, I’ve discovered a reason to deny the representativeness of your sample and a reason to suspect the strength of your inductive generalization.” “Yes, sir.”

At this point, I’m inclined to stop payment on that check.

Now what is supposed to correspond to the tarp-in-the-garden image? Just as the tarp covers a part of the garden that was not examined (because it was too heavy to lift), so too, some portions of abstracta are unsearchable (because—to take one example—we cannot penetrate their complexity given our crude cognitive capacities and tools). Let us inquire, then: Isn’t it likely that if there were a morally justifying reason for permitting the world’s horrific evils, it would be located in a portion of the garden of abstracta that is impenetrable to us?

Several affirmative answers have been presented for evaluation. Here are two for starters: (i) it wouldn’t be at all surprising if it is likely that the magnitude, intensity, duration, and distribution of the evils to be accounted for are themselves so exceedingly complex that a morally justifying reason would exhibit the same feature, a feature that would place it beyond our ken. We are, after all, frequently reminded just how unfathomable is the full and unadulterated history of evil; (ii) it wouldn’t be at all surprising if it is likely that if, owing to a desire to cultivate a certain kind of attitude in his creatures (a desire that would move God if he exists), all the morally justifying reasons are masked by divine intervention and are thus kept safely in obscurity.

Just to be clear—if either of those considerations provides us with a good reason to believe the sample is unrepresentative, the inductive argument for the bridge premise is in real trouble. But there is no need to aim so high and defend such speculative answers; let’s ratchet down—even if we aim lower and claim only that we have no good reason to believe that the sample is representative or that we are in the dark or in doubt about whether the sample is representative, the inductive argument is still in jeopardy. So-called ‘skeptical theists’ are in the business of denying the bridge premise we have been discussing in just this manner. In general, they argue either that we have no good reason to believe (or else that we are

in the dark about whether) the goods we are aware of are representative of the goods that there are (see Wykstra 1984, Alston 1991, Howard-Snyder, 1996, Bergmann, 2001 and 2009, and McBrayer, 2010). And our situation gets even worse.

The Elmer Fudd Problem: Bugs Bunny was a cartoon hero of my childhood. Bugs's nemesis, Elmer Fudd, was forever shotgunning for Bugs but was also astonishingly stupid. All Bugs had to do was throw on a dress and hide his ears, and Elmer immediately mistook him for anything but a rabbit and (disturbingly—for a child's cartoon) usually for prospective love interest. Let us not forget that we have here pretended that (as with rabbits) we know morally justifying reasons when we see 'em. But we should really drop this conceit. We may well be the Elmer Fudds of value theory. Once again, we have no good reason to believe (or else we are in the dark about whether) the entailment relations we know of between goods and permitted evils are representative of the entailments relations there are. Consequently, we may well have already discovered an exceedingly valuable good that would justify God in permitting this or that horrific evil and then (after failing to recognize its necessary connection to the evil in question) mistakenly rejected its candidacy on the grounds that it doesn't require permission of the evil.

And our situation gets even worse, yet again, for we have no good reason to believe (or else we are in the dark about whether) the degree of value we recognize in those goods we are aware of is representative of the total degree of value those goods actually manifest. Consequently, we may well have already discovered a good that would justify God in permitting this or that horrific evil and also have already discovered its necessary connection to the evil in question and then (after failing to recognize its full range of goodness) mistakenly rejected its candidacy on the grounds that it was not sufficiently compensatory.

Consequently, argue our skeptical theists, we have no good reason to think that if there were such a compensating good or some other morally justifying reason for the world's horrific evils, we would be aware of it, or—if it were somehow an object of mere awareness—that we would be able to recognize its full degree of value or its function as a compensating good or morally justifying reason.

Such a realization of our epistemic position does not, of course, amount to a reason to turn theist, but absent any other way of demonstrating the

lack of a compensating good or morally justifying reason, it would (if successful) eliminate many of the most powerful arguments for atheism.

2. For The Bible Tells Me So

A number of religious traditions—including Judaism, Christianity, and Islam—have maintained that among the sources of knowledge available to human beings we should recognize divine revelation. The details differ but in ways that don't matter much for our purposes. I shall focus on the case of Christianity in the following discussion, a religion in which (with a few dissenters here and there) the historical and contemporary view appears to be that God has revealed certain truths touching on matters of consequence to all human persons near and far, past and future, and that these truths are not ones we could have fully discovered left to our own devices (see Swinburne, 2007, and Davis, 2009).

Several models for divine revelation of truths (as opposed to the revelation of God himself) are on offer. Sometimes the proposed mechanism is causal, effected perhaps by dreams, or visions, or some sort of direct neurophysiological tinkering. Sometimes the truth is manifested in some person, some bit of behavior, some miracle, or some other magnificent chain of events. Sometimes the revelation is portrayed as a kind of divine testimony—addressed and spoken to an individual or to a people and communicated directly in God's own voice, or by prophet, or by inspired scripture. (For discussion of the models, see Mavrodes, 1988). However it eventually gets transferred, such testimony has propositional content (Davis, 2009, 35), and the force of the term 'alone' in the phrase "knowledge by revelation alone" is simply once again to signal that human powers of cognition—reason, understanding, imagination, sensation, introspection, memory—are not up to the task of discovering the truth values of these propositions on their own. Apart from revelation (if there is such a thing) we do not have any independent means of verification or any significant epistemic access to the relevant subject matter.

Although particular examples are always contested, candidates for bits of knowledge by revelation alone with respect to the Christian tradition include claims regarding the fall of humanity, the trinity, the incarnation, the atonement, and the general resurrection and life of the world to come. Although specific traditions and creeds aren't really the

focus here and although any doctrine whose credentials are restricted to the testimony of an omniscient and perfectly good being will suffice to generate the conflict I intend to examine, I will need a placeholder in our discussion.

“We look for the resurrection of the dead and the life of the world to come.” So ends the Nicene Creed with a reference to one of the most central teachings of Christianity. Similarly, the Apostles’ and Athanasian Creeds explicitly and prominently call attention to the resurrection of the body. In the discussion to follow I will take this position to imply (at the very least) what I will call *the general resurrection thesis*—the view that every human person who has ever died will rise again from the dead—and in what follows I will invoke this (controversial) thesis as our example of a candidate for knowledge by revelation alone.

3. ‘Notrustem’ Inferences

It is commonly agreed that being deceived is a bad state of affairs and that lying is morally wrong. A more careful pronouncement, however, is that there are some features of being deceived that are of disvalue (e.g. a mismatch between one’s cognitive states and the world) and that lying is *prima facie* morally wrong (i.e. that there is moral reason not to deceive others). Nearly everyone, however, recognizes both that deception can be extrinsically valuable (e.g. it can lead to substantial goods such as saving the deceived from making a life-ruining mistake in a moment of passion) and that the moral presumption against lying can be overridden by even stronger moral reasons in favor of deception on a particular occasion. As is well known, *prima facie* obligations can come into conflict with one another, and when they do, the directive against lying quite obviously is not always the most stringent, overriding, consideration in moral decision-making.

The problem, then, is straightforward. Consider any piece of alleged knowledge by revelation alone, *K*. Moreover—and *this is the surprising part*—simply grant without a fuss that the testimony inviting us to adopt *K* as a truth comes from someone whom we know to be both essentially omniscient and essentially perfectly good. (Of course, traditions may differ on those two aspects of the deity, but I am interested in the problem that emerges on even the most generous list of known divine attributes.)

Here, then, is a first pass at characterizing our puzzling predicament:

Do we thereby have knowledge of K? Well, not if K is false—if K is false, then we are deceived by someone who knows that K is false.

Do we have a way of verifying the truth of K and exposing a deception, if deception it be? Well, not if K is a genuine candidate for knowledge by revelation alone, for earlier we specified we do not have any independent means of verification or any significant epistemic access to the status of such candidates.

Do we have a guarantee that God would not deceive us about whether K is true? Well, not if our being deceived about K is the kind of bad state of affairs for which there exists a compensating good or morally justifying reason. If there is a compensating good or morally justifying reason for such deception, God's essential perfect goodness is not in any way impugned by the deception—on the contrary, it may be morally obligatory to so deceive us.

Do we know that there is no such compensating good or morally justifying reason in this case? There's the rub... apparently not, if we are among those who adhere to the skeptical theist's defensive maneuver for undercutting arguments from evil to atheism. It would seem that consistency would require us to claim ignorance here as before, and for more or less the same reasons, too.

Consequently, our claim to knowledge seems to be threatened: We cannot without reservation trust such divine pronouncements—even if we simply help ourselves to the background assumptions that God exists, that God is essentially omniscient and essentially perfectly good, that God has provided us with his testimony, and that we have interpreted that testimony aright. And once we have lost this particular kind of trust in the testimony, it cannot be the source of testimonial knowledge.

Given the centrality and importance of the sorts of theses identified in Section 2 above as candidates for bits of knowledge by revelation alone, such a 'notrustem' inference points to a severe and underappreciated problem confronting anyone who holds this popular combination of views. (Not entirely underappreciated, however. See Wielenberg, 2010, Maitzen, 2013, and Wilks, 2013.) Finally, it is worth remarking that this problem is significantly different from other critiques of skeptical theism, even those which are explicitly grounded

in claims about moral deliberation and moral knowledge (for such critiques see Almeida and Oppy, 2003, Pereboom, 2005, and Jordan, 2006; for an evaluation of such critiques, see Bergmann and Rea, 2005 and Howard-Snyder, 2010).

In the remaining sections, I will critically examine and evaluate three proposals for responding to this problem.

4. The George Washington Defense

Perhaps we are in the clear on the grounds that God simply cannot tell a lie, since doing so either betrays a kind of weakness or else is always morally wrong, and since God's perfection is incompatible with both defects and moral wrongdoing. After all, the Commandment doesn't say "Thou Shalt Not Lie, Unless, You Know, You Can Wring Some Advantage Out Of It." Moreover, scripture contains a number of passages in which lying and deceiving are subjected to heavy criticism and condemnation. And historical champions such as Augustine, Gregory the Great, and Aquinas were all on board with an absolute prohibition on lying as is too the *Catechism of the Catholic Church* (see Tollefsen and Pruss, 2011).

Such support is frequently hedged, though, and there are considerations that make the prohibition less severe than it seems at first blush. Whereas Kant, in his 1797/1949 essay, "On a Supposed Right to Lie from Altruistic Motives," might actually have advanced the view that "we should let the sky fall before committing the certain sin of telling a lie or practicing deception", it is harder to believe that the sources just cited could have had that view in mind. Four quick comments on this theme:

First (and significantly), the prohibition does not usually extend to all cases of deception but rather to the special case of deception known as lying. As noted above, however, knowledge by revelation could come in many different flavors, only some of which involve the sort of direct assertion that is subject to the charge of lying (as opposed to some other form of deception).

Second, as many authors have noted and sometimes lamented, scripture is also a source of evidence both for the claim that God himself has perpetrated (whether directly or indirectly) an intriguing assortment of lies and cunning deceptions as well as for the claim that such behavior was laudable and morally permissible (see Smith, 2011b).

Third, the *Catechism of the Catholic Church* and the support of certain influential individuals are often qualified by the presence of excepting-clauses of some kind or other. For example, in the case of the *Catechism's* first version, the rule against lying was softened a bit by an accompanying definition of 'lying' which insisted upon "having the right to know the truth" as a necessary condition of being lied to. The prohibition, then, would be silent on cases in which the right had never been present to begin with or else had been either waived or forfeited. In a similar fashion, by imposing restrictions on what counts as a lie, our historical figures are able to consistently advocate an absolute ban on lying while acknowledging the permissibility of some instances of deliberate deception by way of explicit assertion known by the speaker to be false. Compare the fact that Aquinas or the Church could also stand in favor of the prohibition against murder and yet endorse the permissibility of deliberate killings in certain cases of self-defense—not because such actions are permissible murders, but because they do not fall under the scope of the proper analysis of 'murder' at all (see Smith, 2011a).

It would be an odd morality that maintains that deliberate killing in self-defense cases when innocent lives are at stake is acceptable moral behavior, but that lying to achieve the same ends in those very cases would be inexcusable, unqualified moral wrongdoing.

Fourth, it seems dialectically inappropriate to mount a defense of an *ultima facie* duty to refrain from lying by appealing to the very sources whose veracity has been called into question. If I am (in fact) being lied to by Dean Zimmerman (and for my own good), and I come to suspect that this might be the case, I doubt I'll make much useful headway in my attempt to sleuth out the facts by asking Dean if he is lying to me. I'd do better to ask Meghan. If the context in which the question arises of the permissibility of engaging in direct and intentional deception by way of false assertion is one in which scripture and the authority of certain figures is presupposed as reliable and trustworthy ground, then we primarily face the problem of arriving at the correct interpretation of those sources. But that is decidedly not the context in which the question is being raised here.

The George Washington defense against 'notrustem' inferences would presumably have to be founded on independent moral reasoning rather than any appeal to knowledge by revelation alone (or else be subject to a question-begging critique). And when turning to and inspecting that independent moral reasoning, it seems clear that the widespread and

considered opinion remains that despite the disvalue of being misled and the *prima facie* wrongness of lying, the moral presumption against lying can be overridden by even stronger moral reasons in favor of deception on a particular occasion. Once again, then, *prima facie* obligations can come into conflict with one another, and when they do, the directive against lying quite obviously is not always the most stringent, overriding consideration in moral decision making.

5. The Greatest (Tall) Tale Ever Told

Perhaps we are in the clear on the grounds that, whereas God can tell the occasional lie, he simply cannot deceive us about something as important as the fall of humanity, the trinity, the incarnation, the atonement, or the general resurrection thesis.

In other words, the prohibition need not be on the telling of lies as such, but rather on the telling of certain lies rather than others. Well, why? What, exactly, is the problem? Apparently, the idea is that the consequences of being deceived on such momentous topics would be so severe, so unfathomable, so unutterably bad, that despite the intellectual humility ordinarily manifested by the skeptical theist, he can quite clearly discern that an atrocity of this magnitude just could not be justified.

This ploy strikes me as wrongheaded twice over:

First, it is a half-hearted (and vulnerable) skeptical theism which professes ignorance about how much we know about just which things are good, about just how good they are, and about the necessary conditions of their realization in a wide variety of scenarios, but changes its tune in a few special cases to declare that something like *that* just couldn't be tolerated. Whence this confidence? I suspect it might have its origins in the plausible intuition that we know (for example) that there could be no morally justifying reason to permit a world consisting of nothing but sentient beings in devastating pain at every moment of their existence. But even if we do know such a thing, and even if such knowledge is more or less consistent with the general attitudes of skeptical theism elsewhere, note that the extreme scenario before us is not at all relevantly similar to the case of divine deception at issue. A lot gets packed into the 'nothing but' qualification in the phrase "nothing but sentient beings in devastating pain at every moment of their existence." In particular, it has the effect of stipulating a

case in which no other factors (apart from those entailed by that description) are relevant to determining its overall value or justifiability. Perhaps, under the guidance of that ideal stipulation, we are in a special position to see that such a world is inconsistent with God's nature, but there is no guarantee of the absence of other relevant factors in a case of divine deception about the general resurrection thesis, and absent the absence, we have no business temporarily departing from the recommendations of skeptical theism to make confident pronouncements about the intolerability of divine deception on the matters of revelation.

Second, the proposal strikes me as a bit of overexcited hyperbole, a kind of reverse-Panglossianism—"No Worse, There is None / Pitched Past Pitch of Grief." Bad? Yes. But this bad? Suppose that—as several Christians have argued—we human persons are material beings, and suppose further that despite the half-dozen or so proposals for reconciling a materialism for human persons with the doctrine of the general resurrection, this union is simply metaphysically impossible (see Hudson, 2001). Yet perhaps even creatures essentially barred from a certain kind of afterlife are worth creating anyway, and moreover, ought to be deceived on exactly that point. To be clear, that's not my view. I think I am a material object and that materialism for human persons is consistent with the general resurrection thesis, and so I look for the general resurrection and the life of the world to come. But if I'm wrong about those things, and God has deceived me about the prospects for things of my kind, I just can't see that I'm in a position to say that the deception-component of such a state of affairs is as bad as the world-of-unrelenting-and-devastating-pain or anything else that convinces me that it could not be permitted by a morally justifying reason.

6. In God We Trust

My colleague, Michael Rea, has always been a divine friend to me. Of course, that adjective has to be taken in context. Mike isn't God, but he isn't wholly unlike God, either. Mike and God share *some* properties relevant to recognizing him as the fine friend I know him to be. Perhaps Mike and God share some other properties, as well, properties that enable him to impart knowledge to me by way of testimony, despite the fact that I cannot completely trust what he has to say either. It is important

to note, however, that my lack of complete trust need not be based on any failing of Mike's; indeed (as just suggested) it may be predicated on some respect in which Mike resembles God—namely, on his knowledge and goodness.

A brief story: I first gave this paper as an after-dinner talk at a conference organized by Mike at the University of Notre Dame. Mike knows I'm fond of rum and coke, that I tend to get a touch nervous before giving talks, and that if there is rum and coke at hand before I give a talk—it's not at hand for long. Steeling myself against the onslaught of questions soon to be coming from people who had thought more about skeptical theism than I had, I asked Mike if the drink he had just fixed and handed over was indeed a rum and coke. He *knew* the answer. And he had my interests in mind. And he had other goals, too; in particular, he wanted the talk and discussion to go well. He fully understood that my being deceived would have disvalue for me and that his lying would be *prima facie* morally wrong, and yet that he nevertheless may have had other *prima facie* duties in favor of misleading me that were even more stringent in those circumstances. Moreover, as it turns out, on this occasion he did both what he ought to have done and what he knew he ought to have done.

Now it is easy to understand how our story could have concluded in either of two ways: *The first (merely possible) ending:* Mike lied to me (quite justifiably) and handed me a coke. *The second (and actual) ending:* Mike told me the truth (hoped for the best) and handed me a rum and coke. Was my claim to testimonial knowledge that I was then downing a rum and coke imperiled by my concession that I could see how Mike (while knowing the facts of the matter) may have lied to me as a manifestation of his own goodness?

Just to be clear—of course I wouldn't have known if I had in fact been successfully deceived (for then the relevant belief would have been false). The question is rather, given that Mike spoke honestly, did I then have knowledge by way of receiving his (true and known) testimony, despite my realization that there may well have been morally justifying reasons for my being a victim of deception on that occasion? Initially, perhaps, one may think the answer is "of course," and furthermore, "if I can gain testimonial knowledge from Mike in the face of such uncertainty and doubt, surely I can receive it from God."

But how persuasive, on reflection, is transferring that reaction in the Mike-scenario to the case involving God? I have to admit, I'm secretly

hoping it furnishes an adequate response, since I'd like some decent way to reconcile my views and this seems to me the most promising. But allow me to at least state the case for the opposition.

I suspect the problem with this strategy will be found in a feature in which Mike is unlike and inferior to God but which surprisingly makes it all that much easier to acquire knowledge by way of Mike's testimony. Accordingly, let us grant for the sake of argument that I acquired the knowledge that I was drinking a rum and coke when Mike told me so and then explore the difficulties with generalizing on this admission.

Testimonial knowledge is a controversial and hotly-debated topic. The non-reductionists see testimony as a basic source of justification (on a par with perception and memory) and as requiring very little effort from the recipient—it's an innocent-until-given-grounds-for-guilt view; in the absence of relevant defeaters, one may justifiably accept a piece of testimony upon hearing it (see Lackey 2006 and 2011). By contrast, in the (to my mind) more plausible reductionist view, testimony requires something more than the lack of undefeated defeaters to impart justification, and different species of reductionists are divided by their different answers to the 'what-else' question. One popular proposal for the additional ingredient is a non-testimonial good reason for thinking that testimony is generally reliable or trustworthy, while another popular proposal advocates the weaker constraint of a non-testimonial good reason for thinking that a particular instance of testimony is reliable or trustworthy (Lackey, 2006).

Focusing on the weaker version of reductionism (featuring the local rather than global requirement on reliability), I realize that I have a plausible chance of meeting that requirement in my interactions with Mike. After all, he and I are relevantly similar on a wide variety of measures, and I am thereby well-positioned to get good evidence from perception, memory, induction and the like relevant to a judgment of reliability on this occasion. Besides, I'm not altogether a hopeless judge of the likely consequences of my imbibing before presenting, and I share much of Mike's information about whether there are such morally justifying reasons to deceive me in the scenario just related.

In fact, it is precisely because I resemble Mike in a way that neither of us resembles God that I have any prayer of meeting this constraint. But how shall I—with my skeptical-theist hands securely tied behind my back—provide myself with analogous assurance in the case involving God? I am not relevantly similar to the source of revelation. I am not well-positioned

to get good evidence from perception, memory, induction, and the like relevant to a judgment of reliability on this occasion (*more on this crucial point in the following section*). And I am a hopeless judge of what is at stake—cosmically speaking—if I and my fellows are not deceived in some comprehensive, irresistible, and undetectable fashion. Indeed, I am quite utterly in the dark on that matter, as we have been saying.

7. How Bad Can It Get?

So what if the replies are not adequate? As committed skeptical theists could we just take our lumps? Could we learn to live with the possibility of divine deception about the most important of matters? Could we—*wait for it*—could we *bear this false witness*?

Of course we have recourse to the speech with which Descartes once taunted the evil demon: “Let him deceive me about whatever he can, he will never bring it about that I am not, so long as I am conscious that I am.” But, then again, each of us securing his own existence isn’t all that much comfort.

On the assumptions that we do not know there is an absolute moral prohibition on lying and that the characteristic tenets of skeptical theism are true, it is epistemically possible that an essentially omnipotent, essentially omniscient, perfectly good, God deceives us about a *tremendous* number of topics—from whether the world is billions of years old to whether the incarnation occurred, to whether I myself am embodied, to whether I had eggs for breakfast. That should be disconcerting enough; but even worse, on those assumptions, we should apparently be utterly in the dark about whether that is exactly what is happening. Or should we? Can we somehow contain the skeptical threat?

In one case, the answer seems to be clearly in the negative. One unfortunate cost of invoking skeptical theism to combat certain atheistic arguments is to suffer its undermining certain theistic arguments, as well. The fine-tuning argument, for example, contains a premise which asserts that given the fine-tuning thesis, the existence of life-permitting, cosmic conditions is very probable under the hypothesis that God exists (see Collins, 1999). The defense of this premise turns on a line of reasoning linking God’s perfect goodness and omnipotence to the high likelihood of creating a world in which free creatures (like us) can interact with one another.

Why so likely? Because, *everything told*, such a world would be so wonderfully valuable. The skeptical theist, however, is in the business of cautioning us against drawing inferences of exactly that kind. We are simply in the dark, she explains, about whether the existence of creatures (like us) freely interacting with one another ensures a valuable world *all things considered*, for we simply have no idea whether our presence is inconsistent with a much more magnificent good or else requires the permission of something exceedingly evil. Perhaps we should simply concede this unfortunate consequence of skeptical theism (as Bergmann, 2009 does and for just these reasons) and hope that the collateral damage comes to an end there.

After all, why should I think that my perfectly reasonable belief that the cosmos is several billion years old or that I am embodied or that I had eggs for breakfast are even remotely jeopardized? As Michael Bergmann has argued, my knowledge of these truths is not acquired by reflecting on possible goods, possible evils, and the necessary conditions of their realization, but rather in some other—*some independent*—way (Bergmann, 2009). And wasn't it exactly this lack of independent access to verifying truth that made putative cases of knowledge by revelation alone so vulnerable to the commitments of skeptical theism? Surely, then, we can draw a line between the genuine skeptical threats concerning moral matters or divine testimony and the merely apparent skeptical consequences elsewhere.

I'm not so sure; if we get this far, there may be no turning back.

Suppose I temporarily breathe a sigh of relief, reminding myself that many of my beliefs are backed independently by appeal to astronomy, biology, chemistry, geology, physics, intuition, memory, perception, introspection, the natural light of reason, or even Reidian common sense (the last appeal being the focus of Bergmann, 2012). But the relief is not long-lasting. Ian Wilks reintroduces anxiety—*for the theist in particular*—by remarking: “Of course, radical skepticism does not yield so easily. For all we know God has invested us with a delusional tendency of common sense, and on its basis we accept as true, beliefs that are actually false. For all we know there is a greater good served, or evil avoided, by this deception” (Wilks, 2013).

It seems it won't do to simply inquire, “Is my belief that P grounded in some judgment about possible goods, possible evils, and the necessary conditions of their realization?” and then to take skeptical theism

to be a humility-inducing corrective to 'Yes' replies, but altogether silent in the case of 'No' replies. For no matter what alleged independent evidence is marshaled in favor of the thesis at hand, skeptical theism can lead our theist to the view that we should be utterly in the dark about whether investing us with the means to obtain that misleading evidence is the very mechanism that our perfectly-good-and-deceiving creator selected for our arriving at the confident but false judgment that P.

Here's one way of seeing the difficulty: If there is a morally obligating reason for God to deceive me, then I am deceived. If there is no morally justifying reason for God to deceive me, then I am not deceived. If there is a morally justifying reason for God to deceive me, then either I am or am not deceived depending on God's other purposes. Skeptical theists would remind me that I am utterly in the dark about which of those three antecedents is satisfied. And thus the darkness expands so that I am also utterly in the dark about whether I am deceived in the most comprehensive, irresistible, and undetectable fashion.

In short, the skepticism in question threatens to *explode* for our theist—i.e. for anyone who accepts that there is a being possessed of the power to so deceive—quickly moving from a well-contained strategy for opposing a popular style of atheistic reasoning to a near global catastrophe, threatening to undermine the reasonableness of our views in nearly all matters, great and small. (For further discussion of the threats of such explosion and strategies for containment, see Russell, 1996, Beaudoin, 2005, Bergmann, 2009 and 2012, Roeber, 2009, and Wilks, 2009 and 2013).

8. Some Disconcerting Options

So—what to do?

Option I—Skepticism Run Amok: Skeptical theism is non-negotiable, its skepticism spreads uncontrollably (for the theist) to other issues, and we simply have to get used to trading in our confidence in alleged divine revelation for comfort in the thought that even if we are the victims of systematic and far-reaching lies, at least God ain't misbehavin'.

Option II—Revisiting Arguments from Evil: Knowledge by revelation alone is non-negotiable, divine testimony can be and has been bestowed upon us, and we simply have to get used to resisting

arguments from evil in some way other than bemoaning how little we know about goods, their full range of value, and the necessary conditions of their realization.

Option III—Demonstrating Consistency: Skeptical theism is (despite appearances) thoroughly consistent with knowledge by revelation alone. But—*how*—exactly? Here's a hopeful parting thought.

If I were picked up (as if by some unseen hand) and placed in the middle of a labyrinth, I might, after hours or days of unsuccessful attempts to leave, come to the skeptical position that there just might be no way out of its interior. On the other hand, if I found myself unable to escape a maze only after I had wandered in on my own, I would be far less likely to embrace that skeptical hypothesis—after all, if there's a way in, one would think there's a way out.

It seems to me that my current perplexity is more similar to the latter case than to the former. I am unsure which way to turn next, because I accept both skeptical theism and knowledge by revelation alone, and each move toward reconciliation thus far seems to me something of a dead end. Still, I wandered into this mess only after accepting certain theses—that God exists and is both essentially perfectly good and essentially omniscient, that the characteristic tenets of skeptical theism are true, that I have some revelation-based knowledge, that lying is only *prima facie* morally wrong, and so forth. The tension I have been exploring arises only for someone who is inclined, as I am, to take on that particular collection of theses, and so I'm hoping that by unraveling whatever epistemological story successfully explains how I reasonably arrived at that combination of views, I may thereby discover a guide to extricate me from this puzzle.

But I have to say that this particular labyrinth is a dark and uncomfortable place, and I would be delighted if anyone would like to play Ariadne to my Theseus and just get me the hell outta here.¹

¹ For comments and criticism on earlier versions of this paper I thank Michael Bergmann, Godehard Brüntrup, Trent Dougherty, Peter Forrest, John Hawthorne, Daniel Howard-Snyder, Ned Markosian, Trenton Merricks, Sam Newlands, Guo Peng, John Pittard, Michael Rea, Blake Roeber, Jeffrey Russell, Thomas Senor, Amy Seymour, Ge Siyou, Eleonore Stump, Jason Turner, Dennis Whitcomb, and audiences at Logos 2012 (the University of Notre Dame) and at Meiguo Metaphysical Mayhem (Peking University).

References

- Almeida, Michael and Graham Oppy (2003), "Sceptical Theism and Evidential Arguments from Evil", *Australasian Journal of Philosophy* 81, 496–516.
- Alston, William (1991), "The Inductive Argument from Evil and the Human Cognitive Condition", *Philosophical Perspectives* 5, 29–67.
- Beaudoin, John (2005), "Skepticism and the Skeptical Theist", *Faith and Philosophy* 22, 42–56.
- Bergmann, Michael (2001), "Skeptical Theism and Rowe's New Evidential Argument from Evil", *Noûs* 35, 278–96.
- Bergmann, Michael (2009), "Skeptical Theism and the Problem of Evil", in Thomas Flint and Michael Rea (eds.), *The Oxford Handbook of Philosophical Theology* (Oxford: Oxford University Press), 374–99.
- Bergmann, Michael (2012), "Commonsense Skeptical Theism", in Kelly Clark and Michael Rea (eds.), *Science, Religion and Metaphysics: New Essays on the Philosophy of Alvin Plantinga* (Oxford: Oxford University Press), 9–20.
- Bergmann, Michael and Michael Rea (2005), "In Defense of Skeptical Theism", *Australasian Journal of Philosophy* 83, 241–51.
- Collins, Robin (1999), "The Fine-Tuning Design Argument: A Scientific Argument for the Existence of God", in Michael Murray (ed.), *Reason for the Hope Within* (Grand Rapids: Eerdmans), 47–75.
- Davis, Stephen (2009), "Revelation and Inspiration", in Thomas Flint and Michael Rea (eds.), *The Oxford Handbook of Philosophical Theology* (Oxford: Oxford University Press), 30–53.
- Howard-Snyder, Daniel (2010), "Epistemic Humility, Arguments from Evil, and Moral Skepticism", *Oxford Studies in Philosophy of Religion*, 2, 17–57.
- Howard-Snyder, Daniel (1996), "The Argument from Inscrutable Evil", in Daniel Howard-Snyder (ed.), *The Evidential Argument from Evil* (Bloomington: Indiana University Press), 286–310.
- Hudson, Hud (2001), *A Materialist Metaphysics of the Human Person* (Ithaca: Cornell University Press).
- Jordan, Jeff (2006), "Does Skeptical Theism Imply Moral Skepticism?" *Philosophy and Phenomenological Research* 72, 403–17.
- Kant, Immanuel (1797/1949), "Critique of Practical Reason and Other Writings", in Lewis White Beck (ed. and trans.), *Moral Philosophy* (Chicago: University of Chicago Press), 346–50.
- Lackey, Jennifer (2006), "Introduction", in Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony* (Oxford: Oxford University Press), 1–21.
- Lackey, Jennifer (2011), "Testimony: Acquiring Knowledge from Others", in Alvin Goldman and Dennis Whitcomb (eds.), *Social Epistemology: An Anthology* (Oxford: Oxford University Press), 71–91.

- Maitzen, Stephen (2013), "The Moral Skepticism Objection to Skeptical Theism", in Daniel Howard-Snyder and Justin McBrayer (eds.), *A Companion to the Problem of Evil* (Malden: Blackwell), 444–57.
- Mavrodes, George (1988), *Revelation in Religious Belief* (Philadelphia: Temple University Press).
- McBrayer, Justin (2010), "Skeptical Theism", *Philosophy Compass* 5/7, 611–23.
- Pereboom, Derk (2005), "The Problem of Evil", in William E. Mann (ed.), *The Blackwell Guide to Philosophy of Religion* (Malden: Blackwell), 148–70.
- Roeber, D. Blake (2009), "Does the Theist Have an Epistemic Advantage Over the Atheist? Plantinga and Descartes on Theism, Atheism, and Skepticism", *Journal of Philosophical Research* 34, 305–28.
- Russell, Bruce (1996), "Defenseless", in Daniel Howard-Snyder (ed.), *The Evidential Argument from Evil* (Bloomington: Indiana University Press), 193–205.
- Smith, Janet E. (2011a), "Fig Leaves and Falsehoods", *First Things*. Jun/Jul 2011 <<http://www.firstthings.com/article/2011/05/fig-leaves-and-falsehoods/>>
- Smith, Janet E. (2011b), "Why Tollefsen and Pruss are Wrong About Lying", *First Things*. Dec 2011 <<http://www.firstthings.com/onthesquare/2011/12/whytollefsen-and-pruss-are-wrong-about-lying>>
- Tollefsen, Christopher and Alexander Pruss (2011), "The Case Against False Assertions", *First Things*. Sep 2011 <<http://www.firstthings.com/onthesquare/2011/09/the-case-against-false-assertions>>
- Swinburne, Richard (2007), *Revelation: From Metaphor to Analogy* 2nd Edition (Oxford: Oxford University Press).
- Wielenberg, Erik (2010), "Sceptical Theism and Divine Lies", *Religious Studies* 46, 509–23.
- Wilks, Ian (2009), "Skeptical Theism and Empirical Unfalsifiability", *Faith and Philosophy* 26, 64–76.
- Wilks, Ian (2013), "The Global Skepticism Objection", in Daniel Howard-Snyder and Justin McBrayer (eds.), *A Companion to the Problem of Evil* (Malden: Blackwell), 458–67.
- 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.