

1

Evil and Evidence

Matthew A. Benton, John Hawthorne, and Yoav Isaacs

I. INTRODUCTION

The problem of evil presents the most prominent argument against the existence of God. The basic idea of the problem of evil is simple enough for anyone to understand, but there's still some philosophical work to be done regarding it—some general work about epistemology and some specific work about God and evil.

The problem of evil once took the form of an argument claiming a *logical* incompatibility between the existence of God and the existence of evil; such arguments are widely agreed to fail.¹ But soon after, a variety of *evidential* or *probabilistic* arguments arrived. Such arguments typically involve a theological premise, one according to which some sort of evil would be strong evidence against the existence of God,² and an empirical premise, one according to which that sort of evil obtains. Given these two premises (and sometimes other background assumptions), the arguments conclude that God's existence is improbable, perhaps highly improbable.

A prominent response to such arguments goes by the name “skeptical theism.”³ Although authors advocating this approach have taken a wide variety of positions, they share two basic epistemological sensibilities. First, skeptical theists think that we are not in a position to know or competently judge

¹ As acknowledged by Rowe 1979, n. 1, Alston 1991, 29, and many others, due largely to work by Plantinga (esp. 1967, ch. 5, and 1974, ch. 9; see also 2000, 460ff.).

² Unless otherwise noted, we use the term ‘God’ in a fairly loose, minimal way. Our arguments are consistent with, but do not presuppose, the traditional conception of God as omniscient, omnipotent, and omnibenevolent. For these purposes, take the hypothesis that God exists to be the hypothesis that an extremely knowledgeable, extremely powerful, extremely benevolent being exists.

³ This is, as many have pointed out, an inapt label, since one can endorse the skeptical component of skeptical theism without being a theist at all.

which features of the world, if any, could justify God in allowing all the evils of the world. Second, skeptical theists think that this impoverished epistemic position substantially diminishes the force of the evidential problem of evil. Skeptical theists offer a variety of additional claims to substantiate their judgment that our impoverished epistemic position has this effect. This paper will critically scrutinize some of these additional claims. Yet there is a thought—a reasonable thought, a true thought—underlying the core of skeptical theism. This thought can be presented clearly: “I don’t know why God would allow these evils. I certainly wasn’t expecting them. But it’s not crazy to think that God has good reasons for running things this way that I don’t understand. I don’t have a firm grip on the divine mind; God could be up to all sorts of things. So while I agree that it’s strange to think that God would allow these evils, it’s not as strange as some people say.”

But skeptical theists often go on to argue that evil provides no evidence against the existence of God.⁴ They deny that the problem of evil is a problem at all. This is, to our minds, a mistake. Instead, skeptical theists⁵ should deny that the problem of evil is *as much* of a problem as it is often alleged to be.⁶ Even in the absence of a satisfying theodicy, the problem of evil is not a *coup de grâce* against theism. Yet as we’ll argue, evil is clearly evidence against the existence of God, and it can even constitute *a lot* of evidence against the existence of God. Still, anyone who dismisses the possibility that there is some strange, underappreciated reason why God allows for evil makes the problem of evil out to be more problematic than it is.

In §§II–V, we lay out the view of evidence as probability raising and show why evil is evidence against theism, even if it is not evidence against a specific theistic tradition such as Christianity. In §§VI–X we shall consider themes from the skeptical theism literature with which we are dissatisfied (particularly CORNEA, epistemic appearances, radical uncertainty about

⁴ Wykstra (1984, 73, 91) originally claimed that evil was not evidence against theism at all (though he retracted this in Wykstra 1996, 148 n. 14). Others talk this way as well: e.g. Robert Pargetter (1976) and Keith Yandell (1985). Peter van Inwagen (1996, 169–71) says that “While the patterns of suffering we find in the actual world constitute a *difficulty* for theism . . . , they do not—owing to the availability of the defense I have outlined—attain the status of *evidence*” (170–1). Daniel Howard-Snyder and Michael Bergmann (2004, 14) argue for the conclusion that “grounds for belief in God aside, evil does not make belief in atheism more reasonable for us than belief in theism”; Richard Otte argues that “theists should not believe evil, or our ignorance of a good reason for God to permit evil, is evidence against religious belief or the existence of God, at all” (2012, 127), and “at best, the theist should refrain from judgement about whether evil is evidence against the existence of God” (2012, 131); see also Dougherty (2014, §1.2) for discussion.

⁵ Like everyone else.

⁶ Plantinga (2000, 482) seems to strike the right note here. Cf. also Oppy (2013, 50).

prior probabilities, gratuitousness, levering evidence, and the representativeness of goods). Finally, in §§XI–XII, we discuss how evidence works in particular epistemological conceptions of evidence, and consider how easily we might know that there's no God in an atheist world.

II. A PROBABILISTIC VIEW OF EVIDENTIAL STRENGTH

Since we're evaluating arguments about the strength of some evidence, it's worth stating what evidential strength amounts to. We like thinking about evidential strength in probabilistic terms (happily, skeptical theists usually do too).

The standard account of Bayesian confirmation is probability raising—a piece of evidence is evidence for a hypothesis just in case that evidence raises the probability of that hypothesis. (That is, just in case $\Pr(H | E) > \Pr(H)$.)⁷ Similarly, a piece of evidence is evidence against a hypothesis just in case that evidence lowers the probability of that hypothesis. (That is, just in case $\Pr(H | E) < \Pr(H)$.)⁸

In what sort of case will a piece of evidence be evidence for a hypothesis? In what sort of case will $\Pr(H | E) > \Pr(H)$? That's easy—just in case that evidence is likelier to come about if the hypothesis is true than if the hypothesis is false. (That is, just in case $\Pr(E | H) > \Pr(E | \neg H)$.)⁹

These likelihood ratios don't just tell us whether evidence confirms or disconfirms a hypothesis; the likelihood ratios also tell us *how strongly* a piece of evidence confirms or disconfirms a hypothesis.¹⁰ If the evidence is a little bit likelier given the truth of the hypothesis than it is given the falsity of the hypothesis, then it is weak evidence for the hypothesis. If the evidence is much likelier given the truth of the hypothesis than it is given the falsity of the hypothesis, then it is strong evidence for the hypothesis. Specifically, what matters is the ratio of the likelihood ratios, their geometric

⁷ For those unfamiliar with Bayesian epistemology, see Weisberg 2011 for a good overview. For those overly familiar with Bayesian epistemology, we note that we do not intend to weigh in on the subjectivist/objectivist debate about the interpretation of epistemic probabilities.

⁸ Of course, confirmation or disconfirmation only occur relative to an underlying probability distribution. Evidence that confirms a hypothesis in one context can disconfirm that hypothesis in another context.

⁹ This only holds given the assumption that all evidence has nonzero prior probability. We make this assumption in what follows.

¹⁰ Of course, given this standard usage a hypothesis can be strongly confirmed without being probably true. In order to determine whether or not a hypothesis is probably true, both likelihood ratios and prior probabilities are needed.

difference—is the evidence twice as likely if the hypothesis is true? Five times as likely? Ten times? A million times? Any two pieces of evidence with the same ratio of likelihood ratios will have the same effect on a hypothesis. If E_1 is 0.1 likely if H is true and 0.01 likely if H is false and E_2 is 0.05 likely if H is true and 0.005 likely if H is false, then E_1 and E_2 would have the same confirmatory effect on H . E_1 and E_2 have the same ratio of likelihood ratios—10—and so they give equally strong confirmation. The ratio of likelihood ratios is all that matters. Because “the ratio of likelihood ratios” is a bit of a mouthful, confirmation theorists refer to it as “the Bayes factor”. With respect to H , both E_1 and E_2 have a Bayes factor of 10. (We’ll discuss the import of the Bayes factor when we consider “levering” evidence in §IX.)

III. THE PROBLEM OF PARADISE

Many skeptical theists argue that evil isn’t evidence against theism (see n. 4 for some offenders). For example, many skeptical theists seem to endorse one or both of these theses:

No Weight: Considerations pertaining to evil do not disconfirm theism at all.

Non-starter: Evil does not even provide a prima facie reason against theism that would need to be countered by skeptical considerations. (Cf. Dougherty 2014, §1.2)

We’d like to start out by explaining why we think that such theses as *No Weight* and *Non-starter* are deeply misguided.

So, why should one think that evil is evidence against the existence of God? For the same reason anything is evidence against anything—the ratio of likelihood ratios. Intuitively, the probability of there being evil given atheism is higher than the probability of there being evil given theism.¹¹ While it’s a bit hard to say what justifies a particular probability assignment, we can say a bit more about our comparative judgments of the probabilities.

Consider a world of pleasures with no pain, of goods with no evil—an Eden.¹² If the world were like that, then we think that would constitute a fairly overwhelming argument for the existence of God.¹³ In such an Edenic

¹¹ *Contra* Stone (2011, 167): “As the appearance of lots of pointless suffering is as probable given theism as atheism, given CORNEA it cannot lower theism’s probability.”

¹² Hold fixed as best as possible the amount and kinds of goodness of our world, but remove all the evil and suffering.

¹³ Note that our argument does not depend on an Edenic world being fairly overwhelming evidence for the existence of God, but merely on it being some evidence for the existence of God.

world, atheists would face the problem of paradise.¹⁴ But if the probability of God is higher given the complete absence of evil (in an Edenic world), then the presence of evil (as in our world) must reduce the probability of God.¹⁵ Put otherwise: if the absence of evil is evidence for God, then the presence of evil is evidence against the existence of God, and it is misleading for skeptical theists to claim otherwise.

It's a theorem of the probability calculus that:

$$\Pr(H \mid \neg E) > \Pr(H) \leftrightarrow \Pr(\neg H \mid E) > \Pr(\neg H)$$

So if the conditional probability of theism on no evil is greater than the prior probability of theism alone, then the conditional probability of atheism on evil must be greater than the prior probability of atheism alone. Of course, this theorem doesn't settle the *strength* of the bits of evidence. Just because the absence of evil would be overwhelming evidence for God it does not follow that the presence of evil is overwhelming evidence against God. If, for example, the absence of evil were near conclusive proof of God's existence but extremely improbable *a priori*, then the fact of evil would only disconfirm the existence of God a little bit.

(For those who like uses of the principle of indifference, we note an indifference-related result.¹⁶ Let's suppose that the prior probability of God's existence is 0.5 and the prior probability of evil's existence is 0.5. In this case, however much the absence of evil would confirm the existence of God, to that exact same extent the presence of evil must disconfirm the existence of God. If $\Pr(H) = 0.5$ and $\Pr(E) = 0.5$, then $\Pr(H \mid \neg E) = \Pr(\neg H \mid E)$. Given these assumptions, if the absence of evil would be overwhelming evidence for the existence of God, then the presence of evil must be overwhelming evidence against the existence of God.)

Note that one needn't have perfectly precise probabilities in mind to follow this sort of reasoning. There is a wide range of reasonable-seeming probability assignments for which our reasoning holds. The authors of this paper are not in total agreement regarding the prior probabilities, but we all

¹⁴ At least assuming that the existence of atheists is consistent with a world without evil. But one can imagine the epistemic position an atheist would face regardless.

¹⁵ It's been suggested to us that one can avoid this result by claiming that the prior probability that God creates a good world with no evil is zero. It's true that this would avoid the problem, but it leads to untoward results: e.g., were Adam and Eve living in a paradise with such priors, they'd thereby have *no* evidence that there is a God, and in fact would have the strongest possible evidence that there is *not* a God.

¹⁶ Roughly speaking, the principle of indifference mandates that sufficiently comparable hypotheses be given equal prior probability. It's very hard to make the principle of indifference precise without also making it incoherent, so we'll leave our characterization rough.

think that evil is evidence against the existence of God, and do so for similar reasons. If your uncertainty about prior probabilities ranges over probability assignments for which evil is evidence against the existence of God, then—however you resolve your uncertainty about prior probabilities—evil will be evidence against the existence of God.¹⁷

IV. THE PROBLEM OF EVIL AND THE PROBLEM OF IGNORANCE

The world is full of horrors. Pain and death are inflicted in tremendous quantity, and are inflicted both by inexorable nature and by malevolent people. Such horrors could be evidence against the existence of God—this is the problem of evil. We don't know of any good reason why God should allow there to be some of these horrors. Our ignorance of any such good reason could be evidence against the existence of God—this is the problem of ignorance.¹⁸

The problem of evil and the problem of ignorance are related. The problem of ignorance cannot exist without the problem of evil. If there were no evils there could be no problematic mystery about why there are evils. But the problem of evil can exist without the problem of ignorance. Even if our ignorance about evil is not evidence against the existence of God, the evil itself can still be evidence against the existence of God.

We don't think it's unreasonable to be overwhelmed by the problem of ignorance. We see no compelling reason to think that God would reveal to us his reasons for allowing evil. But the weakness of the problem of ignorance does nothing to blunt the problem of evil. Let's divide possible worlds into four categories: [1] theistic worlds with no evils, [2] theistic worlds with evils, [3] atheistic worlds with no evils, and [4] atheistic worlds with evils:

Theism no evil	Theism evil	Atheism no evil	Atheism evil
[1]	[2]	[3]	[4]

¹⁷ We discuss a more radical version of uncertainty about prior probabilities in §VII.

¹⁸ The problem of ignorance fits squarely into the problem of divine hiddenness. We are thinking of the problem of evil narrowly, so that the problem of divine hiddenness is not subsumed by the problem of evil.

Now suppose that one were completely certain *a priori* that one would not know of any good reasons for God to allow evils. Then the fact that one did not know of any good reasons for God to allow evils would have no evidential significance—it wouldn't matter at all. In such a case the existence of evils for which one knows of no good reason would not provide any evidence to think that the actual world is in category [2] rather than in category [4]. Instead, the existence of evils for which one knows of no good reason would just confirm the worlds in categories [2] and [4] at the expense of the worlds in categories [1] and [3], which would be falsified.

It is often presupposed that the problem of evil depends upon it being the case that if there were a God who allowed evils to exist it would be very likely that God's reasons for allowing those evils would be discernible. But this is not so. The problem of evil can proceed without any dependence whatsoever on that sort of expectation. Since we know that there are evils, we know that the worlds in categories [1] and [3] are falsified:

	Theism evil		Atheism evil
[1]	[2]	[3]	[4]

And this is evidence against theism. After all, we argued in the previous section that the theistic worlds are more likely to lack evil than the atheistic worlds—pain and death are a bit surprising in worlds presided over by a benevolent deity. Since evil is more likely given the non-existence of God than given the existence of God, the existence of evil confirms atheism over theism. *A priori* certainty that we would not know of a good reason for God to allow evil changes nothing about the evidential significance of that evil. Our lack of a theodicy may not itself be a problem for theism, but expecting the lack of a theodicy does not do the work that a theodicy would do. The lack of a theodicy may not make things worse for theists, but it certainly doesn't make them better.¹⁹

V. THE PROBLEM OF EVIL FOR RELIGIOUS TRADITIONS

Here is a natural thought: Christianity entails the existence of God, so any evidence against the existence of God is evidence against Christianity.²⁰ This

¹⁹ Otte (2012, 131 and 141–2) makes this mistake.

²⁰ We focus on the case of Christianity, but the same basic pattern of reasoning applies for nearly any religious tradition.

natural thought could hardly be more wrong. Modus tollens is a valid form of deductive argumentation, but there is no probabilistic analogue of modus tollens.²¹ Since Christianity entails the existence of God, the probability of Christianity can never *exceed* the probability of the existence of God—but that’s about it.

Imagine a simple case: There are three distinct possibilities—A, B, and C. A entails $A \vee B$. Now suppose that B is falsified. B’s falsification is clearly evidence against $A \vee B$. But B’s falsification is not thereby evidence against A. Quite the contrary—B’s falsification is evidence *for* A. The probability of $A \vee B$ goes down, but the probability of A goes up.

Even though evil is evidence against the existence of God, it does not follow that evil is evidence against Christianity. In fact, not only is evil not evidence against Christianity, evil is evidence *for* Christianity.²² Christianity entails the existence of evil, so the discovery of evil must confirm Christianity. It may seem strange that the horrors that disconfirm theism confirm Christianity, but it shouldn’t. Consider the horror that an innocent man was unjustly crucified. Such a horror is bad news for theism, but good news²³ for Christianity.

Of course, these formal considerations don’t entail that evil isn’t a problem for Christianity. These formal considerations merely entail that the existence of evil isn’t an *empirical* problem for Christianity. This doesn’t mean that Christianity is off the hook, it just means that whatever problems the existence of evil poses for Christianity must be found in its *prior probability*.²⁴ Observing a friend smiling does not disconfirm the hypothesis that the friend is miserable and smiling. Observing a friend smiling *confirms* the hypothesis that the friend is miserable and smiling. But it’s generally unreasonable to think that smiling people are miserable because the prior probability that someone is smiling and miserable tends to be hugely less than the prior probability that someone is smiling and happy. If it is strange (and it is) that Christian ideology has both a loving, sovereign God and unspeakably

²¹ See Sober 2004.

²² Cf. Rowe’s (1984, 99–100) point about the differential effect of evil on the likelihood of “expanded standard theism” compared with “restricted standard theism.”

²³ We use the term advisedly.

²⁴ This doesn’t mean that there can be no empirical problem for Christianity (or other religious traditions) posed by specific evils. There are evils that are not entailed by Christianity, and thus that might well be evidence against Christianity; whether they are would depend on one’s prior probabilities.

horrific suffering, that strangeness translates into Christianity having a lower prior probability than it would otherwise have.²⁵

We now turn to some specific discussions by skeptical theists.

VI. SEEMINGS AND CORNEA

Early work on the evidential problem centered around epistemic seemings or appearances (Rowe 1979, §2; Wykstra 1984; more recently Matheson 2014 and Tucker 2014). The evidence brought to bear against the existence of God was not evils, but rather what we think about how the evils appear to us. This is, to our minds, a mistake. We are moved by someone crying out, “How could a good God allow such suffering as mine?” We are less moved by someone crying out, “How could a good God allow such epistemic states as mine?”

Rowe’s (1979, 1984, 1996) work on the problem of evil famously considers a fawn painfully dying in a forest fire, and his most succinct version of the argument is this:

(P): No good we know of justifies God in permitting [the fawn’s suffering].

Therefore, it is probable that

(¬G): There is no God. (1996, 270)

Here is a popular way to characterize Rowe’s argument, where (1) and (2) serve to unpack the above (P):

- (1) We can, try as we might, *see no* God-justifying good served by the fawn’s suffering.
- (2) Hence, *it appears that* there *is no* such God-justifying good served by the fawn’s suffering.
- (3) So: *Probably*, there is no God-justifying good served by this suffering.²⁶

²⁵ Calculations of Christianity’s prior probability are problematic, as Christianity’s specificity also makes it less likely *a priori*. It’s hugely less probable that God would incarnate as a man named Jesus than that God would incarnate as a man, but there’s nothing particularly bad about the name “Jesus” so it’s not quite right to hold its attendant improbability against Christianity. If one is methodical one can deal with these issues, but it’s easy enough to make mistakes that we recommend painting in somewhat broad strokes.

²⁶ Wykstra and Perrine 2012, 378; cf. Wykstra 1996, 127.

Most skeptical theists object to the inference from (1) to (2), from claims about what we don't see or know of, to claims about the way things appear.²⁷ Wykstra's "Condition Of Reasonable Epistemic Access," CORNEA, is put forth as a plausible principle which blocks this inference.

(CORNEA) On the basis of cognized situation *S*, human *H* is entitled to claim "it appears that *p*" only if it is reasonable for *H* to believe that, given her cognitive faculties and the use she has made of them, if *p* were not the case *S* would likely be different than it is in some way discernible to her.²⁸

CORNEA has taken many forms.²⁹ Its "if *p* were not the case"-clause is sometimes understood as a sensitivity condition given by a counterfactual conditional³⁰ (thus applying to the closest worlds in which *p* does not hold); at other times it is understood as a probabilistic conditional (thus making a probability space out of all the worlds in which *p* does not hold). Unfortunately, in this case it is unclear whether either form of reasoning is any good at all. Some people apply CORNEA tests holding experiences fixed, and some people apply CORNEA tests letting experiences vary; but arguably it is bad to use CORNEA either way.

VI.1. CORNEA Holding Experiences Fixed

CORNEA doesn't work as advertised when it holds experiences fixed. We don't think that CORNEA is meant to hold experiences fixed. But when people consider worlds like ours in which God exists and has reasons for allowing the evils we see, this is precisely what they do. We thought it would be illustrative to show how disastrous that sort of reasoning is. Consider the following story of Tom and Susan.

²⁷ Note that this notion of appearance cannot be the most normal one. Imagine someone who knows that he is looking at a white box under red lighting. It would be normal for the person to say, "It appears that the box is red, but I know it isn't." But such a claim would violate CORNEA. The person knows that the box isn't red, so the person knows that if the box weren't red it would look exactly as it does. Wykstra and other skeptical theists tend to employ a notion of epistemic appearance first codified by Roderick Chisholm. We don't particularly like this notion, but happily we needn't consider its details in our arguments concerning cornea. See Chisholm 1957, ch. 4, and Wykstra 1984 for more on this notion of epistemic appearance.

²⁸ Wykstra 1984, 85; cf. Wykstra and Perrine 2012, 379.

²⁹ And not just in Wykstra's contributions: see Stone 2011.

³⁰ See the dispute between McBrayer 2009 and Wykstra and Perrine 2012, esp. 384ff.

Tom and Susan are co-owners of a small business. They get along passably well, and not much about them is of note except for the fact that Tom, who keeps the accounting books, has been embezzling money. One day, Susan confronts Tom, shouting, “You jerk! You’ve been stealing from the company!” and proceeds to trash the office by throwing things across the room. Now Tom, being philosophically minded, begins to consider the possibility that Susan’s reasons for doing this relate to his embezzling their business’s funds. (It does seem like a plausible hypothesis.) Tom reasons as follows:

Tom’s Reasoning:

- (4) I can, try as I may, *see no* non-embezzling reason why Susan is trashing the office.
- (5) Hence, *it appears that* there *is no* such non-embezzling reason why Susan is trashing the office.
- (6) So: *Probably*, there is no non-embezzling reason behind Susan’s trashing the office.

But Tom is well-versed in the literature on skeptical theism, and thus thinks to himself, “True, I cannot see any non-embezzling reason why Susan is trashing the office. But does this really support the hypothesis that Susan is doing this because I embezzled our funds? I’m skeptical. My inference to it appearing that there is no non-embezzling reason why Susan is doing this is suspect. After all, in the closest experientially-matching worlds in which Susan has a non-embezzling reason for shouting ‘You jerk! You’ve been stealing from the company!’ and proceeding trash the office, things aren’t likely to be discernibly different. Thus the inference violates *counterfactual CORNEA*. Moreover, in all the experientially-matching worlds in which Susan has a non-embezzling reason for shouting that and proceeding to trash the office, things aren’t likely to be discernibly different. Thus the inference violates *probabilistic CORNEA*. I guess I don’t have reason to think that Susan is trashing the office because I embezzled our company’s money. This is quite a surprising result.”

Tom’s problem (at least his philosophical problem) is that he’s paying attention to unimportant features of his situation, and construing the epistemic significance of those features in a ludicrous way. The important feature of the situation is that Susan shouted out “You jerk! You’ve been stealing from the company!” and proceeded to trash the office—that’s the evidence. It’s *hugely* less likely that Susan would shout these particular words and trash the office conditional on her having a *non-embezzling* reason for her actions than conditional upon her having an embezzling reason for her actions. That’s all that matters. There’s no point at all to thinking about whether it

appears to Tom that Susan had no non-embezzling reason for shouting that and trashing the office.

But it gets worse—since Tom is holding his experience fixed, no proposition that goes beyond his experience can satisfy CORNEA. Suppose Tom's experience doesn't entail that p . The closest $\neg p$ worlds with the same experience as the actual world have the same experience as the actual world. All of the $\neg p$ worlds with the same experience as the actual world have the same experience as the actual world. If one holds experiences fixed, CORNEA is a bad principle.

VI.2. CORNEA Letting Experiences Vary

The intuitive motivation for CORNEA, applied letting experiences vary, is easily grasped:

[S]uppose that your doctor drops a hypodermic needle on the floor, picks it up, looks at it carefully, and proceeds to try to use it on your arm. When you protest that it may be contaminated, he reasons as follows:

- (7) We can, try as we may, see no viruses on the needle.
- (8) Hence, it appears that there are no viruses on the needle.
- (9) So probably (barring defeaters), there are no viruses on the needle.
(Wykstra and Perrine 2012, 379)

Wykstra and Perrine protest that the doctor shouldn't say that it appears that there are no viruses on the needle because

if there *were* viruses on the needle, then given the nature of viruses and human vision, failing to see them is precisely and obviously what you (and the doctor) should expect. For this reason, not seeing such viruses in no way entitles the doctor to claim that there *appear* to be no viruses on the needle. (*ibid.*)

They explicitly intend CORNEA as a formal generalization of this line of reasoning: "CORNEA simply generalizes this intuitive constraint" given by the passage quoted above. Unfortunately, CORNEA can't do what Wykstra and Perrine intend for it to do.

Suppose that the doctor didn't infer that there were no viruses on the needle just from looking at the needle, but instead inferred that there were no viruses on the needle from looking at the needle and rolling a die. Let's say the die landed on 3. Given this evidence—that the needle looks the way it does and that the die landed on 3—the doctor's claim that the needle

appears to be clean easily passes counterfactual CORNEA. It's true that (in a counterfactual sense) *if the needle weren't clean, the situation would likely be discernibly different to the doctor*. In most of the closest worlds in which the needle isn't clean it's not the case both that the needle looks the way it does *and* the die lands on 3—in most of those worlds the die doesn't land on 3. Similarly, it passes probabilistic CORNEA. It's true that (in a probabilistic sense) *if the needle weren't clean, the situation would likely be discernibly different to the doctor*. In most of the worlds in which the needle isn't clean it's not the case both that the needle looks the way it does *and* the die lands on 3—in most of those worlds the die doesn't land on 3. Clearly, it's bad news for CORNEA if all it takes to bypass its constraints is an accessible game of Parcheesi.³¹ And there's a good case to be made that the game isn't necessary: Our evidence is typically fine-grained enough that it's improbable given just about any coarse-grained hypothesis, so CORNEA will be satisfied no matter what.

When CORNEA allows experiences to vary, those very variations mean that it can't do the work it was supposed to do. But is there a principle nearby that can do the work that CORNEA was supposed to do? Is there a principle that can get at what's wrong with the doctor's inference that the needle is clean? Sure there is. The needle is just as likely to look the way it does (die or no) whether or not it's clean, so the way it looks doesn't give any evidence one way or the other. If we want (for some reason) to translate this trivial observation into the language of epistemic appearances, we can tie epistemic appearances directly to likelihoods. Thus we can say that it's wrong for the doctor to think that the needle appears clean (at least if he means "appears clean" as opposed to "appears unclean") because the needle was no likelier to look the way it did conditional upon it being clean than conditional upon it being unclean. Of course, in that case the epistemic appearance does no work and only serves to obscure what's going on in the argument. Having thus tidied up the meaning of (8), the remaining problem is that the inference from (8) to (9) flagrantly commits the base-rate fallacy, for it relies only on the likelihoods of the appearance and ignores the base-rate probability of a needle having a virus on it.³²

³¹ As several authors are aware, counterfactual CORNEA will inherit the problems faced by sensitivity-based accounts of knowledge. We feel this theme is already well understood, and so we won't belabor it.

³² For example, imagine that clean needles always look pristine and that dirty needles usually look sullied but sometimes look pristine. But suppose also that dropped needles are almost invariably dirty. Then although a needle that looked pristine might appear to be clean, it would not probably be clean.

VII. RADICAL UNCERTAINTY ABOUT PRIOR PROBABILITIES

We have argued that evil is evidence against the existence of God. Our arguments are compatible with a moderate degree of uncertainty about the probabilities involved. But our arguments are not compatible with radical uncertainty about the probabilities involved. If one's uncertainty about the prior probabilities for evil and theism led one to entertain all possible prior probability assignments about them or to entertain none at all, then it would be genuinely unclear what import evil had for theistic belief. And many skeptical theists do seem drawn to just this sort of radical uncertainty—but they shouldn't be. Our view is that this radical uncertainty cannot be sensibly maintained, that it can only be (to borrow a phrase from Samuel Johnson) “the last refuge of a scoundrel.”

Skeptical theists are pointedly reserved regarding probabilistic judgments concerning evil and theism. Thus Peter van Inwagen says that “we do not know what to say about the probability of S [the amount, kinds, and distribution of suffering] on theism” (1995, 85). Similarly, Howard-Snyder and Bergmann say,

We just aren't in a position to judge that $\Pr(P \mid G \ \& \ k)$ is low, that it is middling, or that it is high. We should shrug our shoulders and admit that we don't have enough to go on here.³³ (2004, 22)

Note that although skeptical theists claim not to know what to say about these probabilities, they are not worried that the probabilities might be tragically unfavorable for theism. They do not say, “I don't know what sort of evidence evil is concerning God. For all I know, it's tremendous evidence against the existence of God.” Instead they say, “I don't know what sort of evidence evil is concerning God. And I therefore deduce that it isn't any evidence against God at all.” Any such deduction is profoundly dubious.

There are various interpretations of epistemic probabilities. There are broadly subjectivist interpretations, according to which epistemic probabilities only specify an agent's degrees of belief, and there are broadly objectivist interpretations, according to which epistemic probabilities specify the degrees of belief that an agent rationally ought to have. But the claims of the skeptical theists make no sense given either interpretation. If the probabilities at stake are subjective, then there's nothing substantial to

³³ Here G is theism, P is “No good we know of justifies God in permitting E1 and E2 [putatively gratuitous evils]”, and k is our background knowledge.

be ignorant about.³⁴ If the probabilities at stake are objective, then there is something substantial to be ignorant about—and ignorance about it is not functionally equivalent to knowledge that evil is evidentially irrelevant to theism.

Skeptical theists seem to be thinking about probabilities in a non-standard way, in a way that involves both objective and subjective probabilities. Skeptical theists seem to tacitly rely on some sort of epistemological bridge principle connecting objective probabilities and subjective probabilities. Something like

Bridge: If an agent is totally ignorant about the objective evidential significance of p , then p ought to have no subjective evidential significance.

But we can think of no good reason to believe any such principle. Bridge principles connecting beliefs about epistemological significance to epistemological significance are generally problematic, and this one looks worse than most.

Even if our epistemic position regarding God and evil were maximally murky, that murkiness would not dissolve the problem of evil. But we should not judge our epistemic position regarding God and evil to be maximally murky. We have no obviously sound basis for the prior probabilities concerning God and evil, but we have no obviously sound basis for the prior probabilities concerning just about anything else either. In general, we think things through as best we can and believe as seems reasonable to us. We are not entirely comfortable with this blithe approach to epistemology. But we are convinced that the prior probabilities regarding God and evil pose no problems that are not also posed by the prior probabilities regarding cosmology, linguistics, trends in teen dating, the weather, or just about anything else.³⁵ The alternative to our way of thinking is thus not skeptical theism, but total skepticism. At the very least, there is surely no rational basis for maximal skepticism about probabilities concerning God and evil but minimal skepticism about probabilities concerning God and religious experiences, or God and prophetic revelations, or any of the other aspects of philosophical theology about which skeptical

³⁴ An agent might be uncertain about his own levels of confidence, but in that case some quiet reflection might help the agent understand himself better.

³⁵ This lack of obvious foundations is not limited to matters of probability. Our beliefs about logic also lack an obvious foundation. But this needn't be worrisome. In general, one doesn't need to know how one knows something in order to know it.

theists are less than skeptical.³⁶ The special pleading required to get radical uncertainty about the relationship between God and evil, but about nothing else, is untenable.

VIII. TWO KINDS OF “GRATUITOUSNESS” (BOTH GRATUITOUS) AND TWO KINDS OF “POINTLESSNESS” (BOTH POINTLESS)

Many arguments in the skeptical theism literature concern “gratuitous” or “pointless” evils. The existence of gratuitous or pointless evil is meant to be stronger evidence against the existence of God than the bare existence of evil. There are numerous notions of gratuitousness and pointlessness,³⁷ but they may be divided into two basic kinds: one which does not entail that there is no God and one which does entail that there is no God. Neither kind of gratuitousness or pointlessness is of any use, but for different reasons.

Definitions of the first kind tend to be variants of the idea that an evil is gratuitous or pointless if it is not necessary for the existence of some greater good or for the non-existence of some greater evil.³⁸ Such gratuitous or pointless evils are perfectly consistent with the existence of God. Perhaps it’s a good thing for God to give agents libertarian freedom. If those agents were to freely perform evil acts, those evil acts would not be necessary for the good of libertarian freedom. Had those agents performed good acts instead of evil acts the world might well have been strictly better. Endowing creatures with libertarian freedom risks gratuitous evil, but that risk might well be worth taking. Or perhaps there is no maximally good world that God could create, but an infinite progression of better and better worlds. Then whatever world God chose to create would gratuitously lack the goodness of worlds better than it. But that sort of gratuitous lack would be completely unavoidable.

³⁶ Note also that given standard accounts of vagueness in probability assignments, maximal vagueness regarding the relationship between God and evil will, given evil, require maximal vagueness about God.

³⁷ There is now a small literature concerning how to understand “gratuitous” evil: see Rowe 1979, Yandell 1989, Draper 1989, 335 for his notion of “biologically gratuitous,” Alston 1991, 33–4, Rhoda 2010, Hasker 2010, Judisch 2012, 67f., Frances 2013, 16–31, Kraay 2014; cf. also Dougherty 2014, §1.1.

³⁸ A popular variation on this theme defines an evil as gratuitous or pointless if God could have prevented it without thereby preventing some greater good or leading to some greater evil. Definitions of the first kind are not all coextensive, but their differences are irrelevant to our arguments.

Of course, the mere fact that gratuitous or pointless evils are consistent with the existence of God doesn't resolve the problem posed by those evils. It still remains to determine the evidential significance of those evils. But that's always the case; this ideology of gratuitousness and pointlessness does no real work.

Definitions of the second kind tend to be variants of the idea that an evil is gratuitous or pointless if it is such that God could not permit it to exist.³⁹ For clarity, we shall call such evils "strongly gratuitous" or "strongly pointless." Such strongly gratuitous or strongly pointless evils are obviously inconsistent with the existence of God.

The ideology of strong gratuitousness and strong pointlessness accomplishes nothing for the evidential problem of evil. The central question of the evidential problem of evil is this: How much less likely are the evils we see given the existence of God than given the non-existence of God? All parties accept the existence of the evils we see. But—obviously—not all parties accept that the evils we see are strongly gratuitous or strongly pointless. It is therefore tendentious for an argument to assume that the evils we see are strongly gratuitous or strongly pointless.

We also don't like the more modest claim that the evils we see merely *seem* or *appear* to be strongly gratuitous or strongly pointless, for two reasons. First, such an argument is still tendentious. It's doubtful that evils seem to entail God's nonexistence to people who are satisfied by any extant theodicy, or even to people who find it plausible that there is some sound theodicy beyond their ken. Second, these seemings are difficult to characterize precisely, so we'd rather avoid them if at all possible.

Consider the following situation: Suppose that Rachel and Peter are married, but that Rachel has become suspicious that Peter is cheating on her. Rachel is worried because Peter often stays late at his office. This may or may not be reasonable—the evidential import of Peter's habit of staying late depends on many other factors. But it would be obviously tendentious for Rachel to be even more worried because she thinks that Peter is working *philanderingly* late. And it would be silly of her to try to avoid that tendentiousness by being worried that it *seems* to her that Peter is working *philanderingly* late. What Rachel knows is that Peter often works late. She should focus on that, and let the epistemic status of Peter's possible *philandering* fall where it may.

There are two kinds of gratuitous or pointless evils, one which doesn't entail that there is no God and one which does entail that there is no God.

³⁹ A popular variation on this theme defines an evil as gratuitous or pointless if it is not the case that God has sufficient all-things-considered reason to allow it. Definitions of the second kind are not all coextensive, but their differences are irrelevant to our arguments.

Neither kind should be at play when evaluating the evidential problem of evil. The first kind has no impact on evidential reasoning, and the second kind is inappropriate for evidential reasoning. Worse still, the conflation of these two kinds of gratuitousness and pointlessness has engendered much confusion. Since neither sort of gratuitousness or pointlessness is of any epistemological use, the best way to remedy that confusion is to simply dispense with all talk of gratuitous or pointless evils. “Gratuitousness” is gratuitous and “pointlessness” is pointless.

IX. LEVERING EVIDENCE

A recent development concerns whether evil is “levering evidence” against the existence of God, with efforts made at defining what levering evidence is.⁴⁰ Wykstra and Perrine (2012) distinguish three “square” states—square belief, square non-belief, and square disbelief—where square belief corresponds to credences of 0.99 and above, square non-belief to credences of around 0.5, and square disbelief to credences of 0.01 and below. Levering evidence is, for them, evidence which is strong enough to take an agent from one square state to another.⁴¹ To get from 0.01 to 0.5 or from 0.5 to 0.99 you need a Bayes factor of 100 in favor of the hypothesis, and to get from 0.99 to 0.05 or from 0.05 to 0.01 you need a Bayes factor of 100 against the hypothesis. Thus levering evidence must have a Bayes factor of at least 100. On their view, for some evidence to lever a hypothesis that evidence must be at least 100 times as likely to come about if the hypothesis is true than if the hypothesis is false, and for some evidence to lever a hypothesis down that evidence must be at least 100 times as likely to come about if the hypothesis is false than if the hypothesis is true.

But talk of square states is liable to cause confusion. Suppose an agent’s credence starts at 0.011 and on the basis of some evidence his credence goes to 0.989—his credence didn’t go from one square state to another; he went from just above the bottom square state to just below the top square state. But that change took more evidence than it would take to go from one square state to an adjacent square state. Worse still, it looks like no evidence

⁴⁰ Wykstra (1996, 130ff.) initially coined the phrase in interpreting what he called “Rowe’s Weightiness Thesis” (see Rowe 1979, 337–8). Wykstra 2007 used the Carnapian language of “dynamic support”; but in Wykstra and Perrine 2012, 380–1, in response to McBrayer 2009, “levering” is used again.

⁴¹ Wykstra and Perrine (2012, 381) call such a doxastic change “sharp.”

could lever up a proposition with 0.6 probability and similarly that no evidence could lever down a proposition with 0.4 probability.

To avoid these sorts of problems, Wykstra and Perrine refine the idea of levering evidence:

[W]e may think of ourselves like Blind Lady Justice. *Eventually* she wants to be in a position to weigh all the evidence for and against a particular hypothesis in the two pans of her scales. But *initially*, she wants to determine how weighty some individual pieces of evidence are. She thus starts from some squarely neutral or “even” position, and asks how much some item of interest can tilt her balance from that even position. (2012, 382)

To avoid problems caused by variations in the prior probabilities of propositions, Wykstra and Perrine ask us to hypothetically place each proposition in question into the same position—square non-belief. And how do we get our propositions into this middling state? It cannot just be by imagining away evidence.⁴² After all, one might thereby imagine away evidence that affects the significance of the evidence in question. Imagine flipping two fair coins and seeing that the first coin landed heads. Then if you see that the second landed heads you know that the two coins landed the same. But without that evidence the fact that second coin landed heads isn't evidence that both coins landed the same. For similar reasons you can't just add evidence to get to 0.5—lots of different collections of evidence could get you to 0.5 and the significance of the evidence in question is liable to vary. There's only one way to make their proposal work: hold all of your credences conditional upon H and all of your credences conditional upon $\neg H$ fixed, but shift the weights of your credences around so that your credence in H is 0.5. That is, Jeffrey conditionalize over partition $\{H, \neg H\}$ and values $\Pr(H) = 0.5$, $\Pr(\neg H) = 0.5$. Having done that, see if conditionalizing on the evidence in question gets you to either 0.99 or 0.01. If it does, it's levering evidence; if it doesn't, it isn't. And when will that conditionalization get you to either 0.99 or 0.01? Just when the Bayes factor of the evidence is at least 100.⁴³

⁴² We certainly don't think that 0.5 is a default probability for when there's no evidence. But even if one has more sympathy for the principle of indifference (as Wykstra seems to) it's of no help in this case.

⁴³ Wykstra and Perrine propose an epistemological principle which they call “core”: “In cognitive situation S giving new input E , E is levering evidence only if it is the case that: (*crux*) if H were false, then E would likely be different” (2012, 377). This principle is quite true; in fact, it is obviously true—though their exposition of it can obscure this. To clarify this principle, Wykstra and Perrine give it a probabilistic interpretation and dub it P-CORE: “(P-CORE) In cognitive situation S giving new input E , E is levering evidence only if it is the case that: (*p-crux*) the conditional probability of E on not- H —*viz.*, $P(E \mid \text{not-}H \ \& \ k)$ —is below 0.5”. (2012, 392; ‘ k ’ is one's background knowledge.)

The underlying idea of levering evidence makes sense; it's a simple idea. Levering evidence is just rather strong evidence. But the machinery developed to define the notion of levering evidence does no real work. Still, a few facts about levering are worth noting, if only to forestall confusion.

Even if e levers p and p entails q , it does not follow that e levers q . Any proposition entails any tautology, and no evidence can lever a tautology—tautologies have probability of 1, and they don't budge. The transitivity of levering fails spectacularly even with contingent propositions. Suppose that a perfect detective is going to investigate a suspect for a crime. This perfect detective will either discover conclusive proof of the suspect's innocence or conclusive proof of the suspect's guilt. Whichever sort of conclusive proof he finds, he'll place it in one box randomly selected from one hundred. The detective will then reveal the number of the box with the proof. When the detective calls out (say), "34!" he levers both the hypothesis that conclusive proof of guilt is in box 34, call it G (which entails that the suspect is guilty), and the hypothesis that conclusive proof of innocence is in box 34, call it I (which entails that the subject is innocent). It levers G and I because each of these hypotheses had a prior probability of 0.005, that is, square disbelief; and after hearing the call of "34!" each credence is levered up to 0.5, square non-belief. But the call of "34!" has no bearing whatsoever on whether the suspect is guilty or innocent; it just means that the randomly selected box was numbered 34. In general, transitivity fails for levering. Thus it's possible for some evidence to lever a proposition that entails God's nonexistence without levering God's nonexistence.

X. THE REPRESENTATIVENESS OF GOODS

Skeptical theists are willing to grant that no goods we know of justify horrendous evils. But skeptical theists deny that this fact licenses inferring that no good whatsoever justifies horrendous evils. After all, the thought goes, perhaps there are goods beyond our ken which justify God in allowing

P-CORE only states that evidence can't lever a hypothesis unless it's unlikely given the falsity of that hypothesis. But as we've seen, to lever a hypothesis some evidence has to be at least 100 times as likely given the truth of that hypothesis than it is given its falsity. But the likeliest any evidence can be given a hypothesis is certain—the likeliest anything can be is probability 1. Therefore the likeliest any levering evidence for a hypothesis could be given the falsity of that hypothesis is 0.01. It's obvious that evidence can't lever a hypothesis if it's not particularly unlikely given the falsity of that hypothesis. Levering evidence is strong evidence. To be strong evidence for a hypothesis that evidence has to be much likelier if the hypothesis is true than if the hypothesis is false. And if the evidence is much likelier if the hypothesis is true than if the hypothesis is false, then the evidence has to be pretty darn unlikely if the hypothesis is false.

horrendous evils. Michael Bergmann defends just this line of thought. Unfortunately his way of defending this line of thought is deeply flawed. Bergmann frames his view in terms of skeptical theses about representativeness. For example, Skeptical Thesis 1:

ST1: We have no good reason for thinking that the possible goods we know of are representative of the possible goods there are. (Bergmann 2001, 279; 2009, 376)

Unfortunately, reflection about whether or not the possible goods we know of are representative of the possible goods there are can grant no purchase on whether unknown goods justify God in allowing the evils we see.

Bergmann correctly notes that representativeness facts are property-relative:

[A] sample of *x*s can be representative of all *x*s relative to one property but not another. For example, a sample of humans can be representative of all humans relative to the property of *having a lung* while at the same time not being representative of all humans relative to the property of *being a Russian*. (2009, 377)

However, this fact dooms thought about representativeness to uselessness.

It's obvious that the possible goods we know about will be representative of the possible goods there are in some respects. For example, they will all be goods. It's obvious that the possible goods we know about will not be representative of the possible goods there are in some respects.⁴⁴ For example, the possible goods we know about will be unlike the possible goods there are in that we know about all the possible goods we know about whereas we do not know about all the possible goods there are.

So, which property-relative representativeness is relevant to evaluating the problem of evil? As far as we can tell, the relevant property is *failing to justify God for allowing horrendous evils* (Bergmann agrees: 2012, 12). And *ex hypothesi* (for the skeptical theist) none of the possible goods we know about justify God for allowing horrendous evils. So, in that respect, are the possible goods we know about representative of the possible goods there are? Well, that depends. If no possible goods justify God in allowing horrendous evils, then the goods we know about are representative. If some possible goods justify God in allowing horrendous evils, then the goods we know about aren't representative. So the question of whether or not the possible goods we know about are representative of the possible goods there are amounts to the question of whether or not any unknown good justifies God

⁴⁴ At least if there are any possible goods we don't know about. If we do know about all the possible goods there are, then, trivially, the possible goods we know about will be representative of the possible goods there are relative to all properties.

in allowing horrendous evils. But the question of whether any unknown good justifies God in allowing tremendous evils is what motivated our investigations into representativeness in the first place! Bergmann's talk of the representativeness of goods amounts not to a solution but to a repackaging of the original problem.

There is another way to state the difficulty with Bergmann's approach. Bergmann claims that his skeptical theses (see Bergmann 2012, 28) seem highly plausible to both theists and non-theists alike, and we've been focusing on ST1. ST1 states that we have no good reason for thinking that the goods we know about are representative of the possible goods there are. Recall that the property of interest is that of failing to justify God for allowing horrendous evils, and the attempted point of this reflection on representativeness was to give some purchase on whether or not there are any God-justifying reasons. So who is this "we" to which Bergmann alludes? The "we" of ST1 cannot include theists who think they have an adequate theodicy to explain the existence of evil. Once you believe that we know of a good which justifies God in allowing evil, the question of whether unknown goods contain the same relative frequency of God-justifying goods becomes entirely otiose. The "we" of ST1 also cannot include atheists who think that the evidential problem of evil gives strong reason to disbelieve that there is a God. By believing that the evidential problem of evil is serious such atheists must both believe that they don't know of any good which defuses the problem and that it is not plausible that an unknown good defuses the problem. So who is the "we" of ST1? Having excluded non-skeptical theists and non-skeptical atheists, it must be populated by various shades of skeptics. But ST1 is supposed to provide an argument *for* Bergmann's skeptical stance; it shouldn't simply be Bergmann's skeptical stance. But it is.

XI. THE PHENOMENAL CONCEPTION OF EVIDENCE

We said early on that evil is evidence against the existence of God, and we've evaluated evidence in a probabilistic framework. But we haven't yet said what evidence actually is. There are two notions of evidence that we like:⁴⁵ the phenomenal conception of evidence and the knowledge-first conception of evidence. We will treat the former in this section, and the latter in the subsequent section (§XII). It's worth exploring these notions in detail, as each notion raises some interesting issues for the problem of evil.

⁴⁵ Which is to say that there is a bit of friendly disagreement among the authors.

According to the phenomenal conception of evidence, an agent's evidence is his phenomenal state.⁴⁶ Experience—and only experience—is evidence. The phenomenal conception of evidence has a broadly internalist feel, and (of course) myriad arguments pro and con.

Given the phenomenal conception of evidence, we hope your evidence right now does not include any horrors. We presume that you are not currently suffering tremendous pain. But if you're not currently suffering tremendous pain then you can't be certain that any such tremendous pain exists. You may think it plausible that there is tremendous suffering (we confess that we find this plausible too) but that there is tremendous suffering is simply not part of your evidence. Thus the phenomenal conception of evidence allows for what we like to call *really* skeptical theism—theism that denies that the horrors that motivate the problem of evil exist. Why should we believe that fawns suffer in forest fires? Perhaps God graces fawns with inner tranquility while their bodies writhe in flame. Perhaps there are no fawns in forest fires at all, and the inferences we make that make us believe in dying fawns are simply wrong. We think that really skeptical theism is a bit silly, but that it is silly is important. Even conditional on the existence of God, it seems wrong to doubt the existence of terrible suffering. But that rather makes it seem that suffering may not be so obviously pointless as is sometimes suggested. If it is much less strange to think that God has reasons to allow for horrendous evils than to think that God has reasons to allow for the misleading impression that there are horrendous evils, then thinking that God has reasons to allow for horrendous evils must be far from maximally strange. Skeptical theism is significantly more credible than really skeptical theism. That may not be much, but it is something.

XII. KNOWLEDGE AS EVIDENCE: E=K

According to the knowledge-first conception of evidence, an agent's evidence is what she knows (E=K): the set of the propositions an agent knows is that agent's evidence. Experiences can *provide* evidence, but that's because experience can yield knowledge; acquiring new evidence *is* acquiring new

⁴⁶ Plausibly, it's better to think of the agent's evidence as the proposition that fully describes the intrinsic character of his phenomenal state. The phenomenal conception of evidence has various possible precisifications, but such technicalities do not concern us here.

knowledge⁴⁷ (Williamson 2000, 196–7; 2014, 4). The knowledge-first conception of evidence has a broadly externalist feel, and (of course) myriad arguments pro and con. Knowledge-first epistemology has distinctive implications for agents in theistic worlds who know that theism is true and for agents in atheistic worlds who know that atheism is true. We treat these in turn.

XII.1. Knowledge of Theism

Imagine an agent who knows that theism is true. If, when witnessing evil, that individual only gains new evidence, that agent will never acquire a body of evidence on which theism is unlikely. After all, the likelihood of theism on any set of propositions that includes theism is 1. In this case the prior probability of theism (and any attendant uncertainty about the prior probability of theism) would be immaterial. In order for facts about evil to pose a deep epistemic problem for such an individual, such facts would have to somehow make certain evidence—in particular knowledge of theism—go away. The question would be whether facts about evil defeat knowledge of theism. (This sort of defeat shouldn't be modeled probabilistically; it's very hard to contrive probabilistic models of why certain evidence disappears.) It is not easy at all to see why evil would force theistic evidence to go away. This kind of anti-theistic project would require a persuasive model of defeat for such cases, and we know of no such model.^{48, 49}

The knowledge-first framework allows for a kind of theism that is very concessionary to evidential arguments from evil. This sort of theist concedes that, conditional on facts about evil alone, theism is unlikely. But this theist then reminds us that an agent's probabilities depend on that agent's entire body of evidence. The theist we have in mind contends that his or her own evidence consists not only in various facts of evil but also in a body of theism-entailing propositions. And quite obviously, theism would not be unlikely on *that* body of evidence. Call this 'Concessionary Theism'. The

⁴⁷ Of course, some of the known propositions will be about the agent's experiences. But those propositions needn't all be (and generally won't all be) confined to facts about the agent's experiences.

⁴⁸ Of course, facts of evil might in certain cases induce a loss of belief and ipso facto induce a loss of knowledge. But that hardly constitutes a general model of defeat. For further discussion of defeat see Maria Lasonen-Aarnio (2010, 2014), and Max Baker-Hyatt and Matthew A. Benton (forthcoming). Even if you're more open to defeat than these authors, it's far from clear that evil would defeat theistic knowledge.

⁴⁹ Relatedly, one might think that the knowledge of evil blocks the acquisition of theistic knowledge. Again, this kind of anti-theistic project would require a persuasive model of blocking for such cases, and we know of no such model.

concessionary theist can find all sorts of common ground with the typical proponent of the evidential argument from evil. She can agree that theism is unlikely on the evidence currently possessed by such a proponent. She can fully understand why, for all the proponent knows, theism is unlikely on the evidence that the proponent thinks is available to her, the theist. After all, the proponent does not know that the theist's evidence includes the fact that God exists. (On the E=K model, any occasion when we don't know whether or not a person knows a certain proposition is a case where we don't know what that person's evidence is.) Nevertheless, the theist can explain to the proponent why she is unmoved by the problem of evil: Theism would hardly be unlikely on a body of evidence that included theism. Further the theist will retain hope that the proponent will one day join her in possessing an enhanced body of evidence.⁵⁰ Obviously, the E=K approach does not mandate Concessionary Theism. But at the very least Concessionary Theism shows how seriously a theist can take the problem of evil without thereby endangering her theistic convictions.

Absent a compelling model of defeat, theists have nothing much to fear from arguments from evil. Of course an E=K approach will hardly encourage atheists to think theists possess theism as evidence. But it may encourage them to concede that if theists at one time know theism, then learning facts about evil does not make such knowledge go away.

XII.2. Knowledge of Atheism

How easy is it to know propositions that entail the falsity of theism? Our view is that—so long as theism is, in fact, false—it's surprisingly easy.

Supposing that theism is false, imagine that a person witnesses a vicious mugging of her friend, and forms the belief:

- (10) A morally praiseworthy person would have prevented a mugging of my friend at this time if s/he had the opportunity.

Can this belief constitute knowledge? Here is one reaction: "Obviously the person cannot know (10). There could be a morally praiseworthy person who wouldn't stop the event given the opportunity because they didn't realize how bad muggings are. Or maybe the morally praiseworthy person

⁵⁰ Note that, in general, possessing at time t evidence on which a hypothesis H is unlikely hardly constitutes a general block on later possessing H as evidence. An agent might possess at t evidence on which it is unlikely that he will have a headache at $t+1$. But he might still at $t+1$ acquire the proposition that he has a headache at $t+1$ as evidence.

foresaw that something even worse would happen if the mugging were prevented. So at the very least, (10) has to be qualified.” From the E=K perspective, this reaction is wrongheaded. Not all possibilities need be reckoned with; some possibilities are epistemically irrelevant.⁵¹ (10) could certainly be true as it stands in some atheistic worlds. It might well be that all of the close possible worlds where a morally praiseworthy person has the opportunity to prevent the friend’s mugging are ones where the mugging is prevented. And a very plausible case can be made that a person making judgments like (10) can know (10) so long as cases where such judgments are false are sufficiently dissimilar to the case in which the judgment is made.⁵² Remote possibilities are precisely those that tend to be epistemically irrelevant. The existence of worlds where judgments like (10) go wrong is no more damning to the possibility of knowing (10) than the existence of brain-in-a-vat worlds is damning to the possibility of knowing that one has hands.

Now supposing that (10) can be acquired as evidence in non-farfetched types of atheist worlds,⁵³ in such worlds an atheist can acquire evidence that entails that there is no eternally existing omniscient, omnipotent, and morally praiseworthy being. Given that being morally perfect entails being morally praiseworthy, such evidence will also entail that there is no eternally existing omniscient, omnipotent, and morally perfect being.⁵⁴ Assuming a standard counterfactual logic, propositions of the forms

- (11) If there were an *F* object, there would not be a *G* event
- (12) There is a *G* event

together entail

- (13) There is no *F* object.

⁵¹ The character of the agent’s environment will play a role in determining which worlds are epistemically irrelevant. The character of this role is one of the central issues in philosophical disputes about knowledge.

⁵² And supposing that the practice of making judgments such as (10) would not easily lead the person to fall into error.

⁵³ Note that the acquisition of such evidence does not require that one be able to provide compelling arguments against hypotheses compatible with the superficial evidence and incompatible with (10). And as before, the acquisition of such knowledge does not require that the agent would not believe (10) if the world were such that a morally praiseworthy being had the opportunity to prevent the mugging but did not do so for some hidden reason.

⁵⁴ We note that any evidence that entails that there is no God (including the proposition that there is no God) is levering evidence against the existence of God no matter what.

Hence the truth of (10) entails there is no morally praiseworthy being that had the opportunity to prevent the mugging. But the proposition that there is an omniscient, omnipotent, morally praiseworthy being entails that there is a morally praiseworthy being that had the opportunity to prevent the mugging. So it seems that in appropriate sorts of atheistic worlds, someone can acquire evidence that entails that standard theism is false by making an offhand judgment about a mugging—no explicit philosophy of religion required!⁵⁵

A few additional observations. First, note that from an E=K perspective, the standard distinction between evidential and logical arguments from evil is potentially misguided. We have just seen that from that perspective it is coherent to allow that we acquire evidence that *entails* that theism is false. In this case, again, the prior probability of theism (and any attendant uncertainty about the prior probability of theism) would be immaterial. Second, note that even supposing that (10) can be acquired as evidence, there is a way of blocking the conclusion that theism is knowably false in the relevant worlds. One can deny epistemic closure, that is, one can deny that agents are always in a position to know the entailments of what they know.⁵⁶ That said, closure denial is a minority view and we do not think this a very plausible way of blocking atheistic knowledge in the relevant worlds. Third, even supposing atheistic knowledge is possible via the means we've indicated, it hardly follows that the atheist will have a dialectically effective way of making progress against the theist. After all, the theist will believe that the actual world is not epistemologically hospitable for atheism. Even supposing the world is epistemologically hospitable for atheism, the theist will not know that it is and will not know that the atheist knows such atheism-entailing propositions as (10). Hence it will be tendentious and unprofitable for the atheist to attempt to make dialectical progress by deploying (10) as a premise. Still, the question whether the atheist can gain the dialectical upper hand is a sharply different question from whether the atheist can know the truth of atheism on the basis of evil. And from a plausible E=K perspective, typical skeptical theist misgivings are almost entirely beside the point.

⁵⁵ There are numerous such cases. Someone sitting might know that if there were a tiger in front of him, he'd run. This would entail the nonexistence of an invisible, omnipresent tiger. Someone unloved might know that he'd be happier if someone loved him. This would entail the nonexistence of an all-loving God. And so on. For related discussion, see Billy Dunaway and John Hawthorne forthcoming.

⁵⁶ See Hawthorne 2004, 34ff. for some such closure theses.

In the previous section we described one way that the theist might be concessionary: It is perfectly coherent for a theist to think that theism is unlikely on a body of evidence that is restricted to mundane facts about evil. In this section we have in effect outlined a further way that the theist might be concessionary: She may concede that were the world as the atheist thinks it is, exposure to evil might provide the atheist with evidence that is incompatible with God. These concessions are of a different order. The first is a concession about the probabilistic implications of what evidence the theist thinks the atheist is actually in possession of. The second is a concession about what evidence is accessible to the atheist in situations that the theist thinks are non-actual. But as far as we're concerned both kinds of concessions are ones that theists should seriously consider.

XIII. CONCLUSION

In this paper we've expressed dissatisfaction with many of the avenues down which skeptical theists have traveled. We've argued that appeals to CORNEA, to epistemic appearances, to gratuitous evils, to levering evidence, and to the representativeness of goods accomplish very little as responses to the evidential problem of evil. Following our advice might make the project of theodicy all the more attractive; but doing so need not lead the skeptical theist to despair, for we've shown how much of the discussion can turn on the specifics of one's theistic tradition, and on what notion of evidence one employs. Not all is doom and gloom. In particular, we've shown that given either the phenomenal conception of evidence or E=K, the skeptical theist can gain new responses to the problem of evil.⁵⁷

REFERENCES

- Alston, William P. 1991. "The Inductive Argument from Evil and the Human Cognitive Condition." In James E. Tomberlin (ed.), *Philosophical Perspectives*, vol. 5, 29–67. Atascadero: Ridgeview Publishing.

⁵⁷ For helpful feedback, we're grateful to Charity Anderson, Trent Dougherty, Branden Fitelson, Maria Lasonen-Aarnio, Richard Swinburne, Vince Vitale, and William Wood, and to workshop audiences at Rutgers University and the University of St. Thomas. This publication was made possible by support from ANU, Oxford, and through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

- Baker-Hyatt, Max and Benton, Matthew A. forthcoming. "Defeatism Defeated." *Philosophical Perspectives*, vol. 29.
- 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 P. Flint and Michael Rea (eds.), *The Oxford Handbook of Philosophical Theology*, 374–400. Oxford: Oxford University Press.
- Bergmann, Michael. 2012. "Commonsense Skeptical Theism." In Kelly James Clark and Michael Rea (eds.), *Reason, Metaphysics, and Mind: New Essays on the Philosophy of Alvin Plantinga*, 9–30. Oxford University Press.
- Chisholm, Roderick M. 1957. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Dougherty, Trent. 2014. "Skeptical Theism." In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/spr2014/entries/skeptical-theism/> Spring 2014 edition.
- Draper, Paul. 1989. "Pain and Pleasure: An Evidential Problem for Theists." *Noûs* 23: 331–50.
- Dunaway, Billy and Hawthorne, John. Forthcoming. "Scepticism." In William J. Abraham and Frederick D. Aquino (eds.), *The Oxford Handbook of the Epistemology of Theology*. Oxford: Oxford University Press.
- Frances, Bryan. 2013. *Gratuitous Suffering and the Problem of Evil: A Comprehensive Introduction*. New York: Routledge.
- Hasker, William. 2010. "Defining 'Gratuitous Evil': A Response to Rhoda." *Religious Studies* 46: 303–9.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Howard-Snyder, Daniel and Bergmann, Michael. 2004. "Evil Does Not Make Atheism More Reasonable than Theism." In Michael L. Peterson and Raymond J. VanArragon (eds.), *Contemporary Debates in Philosophy of Religion*, 13–24. Malden: Blackwell Publishers.
- Judisch, Neal. 2012. "Meticulous Providence and Gratuitous Evil." In Jonathan L. Kvanvig (ed.), *Oxford Studies in Philosophy of Religion*, volume 4, 65–83. Oxford: Oxford University Press.
- Kraay, Klaas J. 2014. "Peter van Inwagen on Gratuitous Evil." *Religious Studies* 50: 217–34.
- Lasonen-Aarnio, Maria. 2010. "Unreasonable Knowledge." *Philosophical Perspectives* 24: 1–21.
- Lasonen-Aarnio, Maria. 2014. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88: 214–345.
- McBrayer, Justin P. 2009. "CORNEA and Inductive Evidence." *Faith and Philosophy* 26: 77–83.
- Matheson, Jonathan D. 2014. "Phenomenal Conservatism and Skeptical Theism." In Trent Dougherty and Justin P. McBrayer (eds.), *Skeptical Theism: New Essays*, 3–20. Oxford: Oxford University Press.

- Oppy, Graham. 2013. "Rowe's Evidential Arguments from Evil." In Daniel Howard-Snyder and Justin P. McBrayer (eds.), *The Blackwell Companion to the Problem of Evil*, 49–66. Malden: Blackwell Publishing.
- Otte, Richard. 2012. "Comparative Confirmation and the Problem of Evil." In Jake Chandler and Victoria S. Harrison (eds.), *Probability in the Philosophy of Religion*, 127–43. Oxford: Oxford University Press.
- Pargetter, Robert. 1976. "Evil as Evidence against the Existence of God." *Mind* 85: 242–5.
- Plantinga, Alvin. 1967. *God and Other Minds: A Study of the Rational Justification of Belief in God*. Ithaca: Cornell University Press.
- Plantinga, Alvin. 1974. *The Nature of Necessity*. Oxford: Clarendon Press.
- Plantinga, Alvin. 2000. *Warranted Christian Belief*. New York: Oxford University Press.
- Rhoda, Alan R. 2010. "Gratuitous Evil and Divine Providence." *Religious Studies* 46: 281–302.
- Rowe, William L. 1979. "The Problem of Evil and Some Varieties of Atheism." *American Philosophical Quarterly* 16: 335–41.
- Rowe, William L. 1984. "Evil and the Theistic Hypothesis: A Response to Wykstra." *International Journal for Philosophy of Religion* 16: 95–100.
- Rowe, William L. 1996. "The Evidential Argument from Evil: A Second Look." In Daniel Howard-Snyder (ed.), *The Evidential Argument from Evil*, 262–85. Bloomington: Indiana University Press.
- Sober, Elliott. 2004. "The Design Argument." In William Mann (ed.), *The Blackwell Guide to the Philosophy of Religion*, 117–47. Malden: Blackwell Publishers.
- Stone, Jim. 2011. "CORNEA, Scepticism, and Evil." *Australasian Journal of Philosophy* 89: 59–70.
- Tucker, Chris. 2014. "Why Skeptical Theism Isn't Skeptical Enough." In Trent Dougherty and Justin P. McBrayer (eds.), *Skeptical Theism: New Essays*, 45–62. Oxford: Oxford University Press.
- van Inwagen, Peter. 1995. *God, Knowledge, and Mystery: Essays in Philosophical Theology*. Ithaca: Cornell University Press.
- van Inwagen, Peter. 1996. "The Problem of Evil, the Problem of Air, and the Problem of Silence." In Daniel Howard-Snyder (ed.), *The Evidential Argument from Evil*, 151–74. Bloomington: Indiana University Press.
- Weisberg, Jonathan. 2011. "Varieties of Bayesianism." In Dov M. Gabbay, Stephen Hartmann, and John Woods (eds.), *Handbook of the History of Logic: Inductive Logic*, volume 10, 477–51. Oxford: North Holland/Elsevier.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy. 2014. "Knowledge First." In Matthias Steup, John Turri, and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 2nd edition, 1–9, 22–5. Malden: Wiley.

- Wykstra, Stephen J. 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 J. 1996. "Rowe's Noseeum Arguments from Evil." In Daniel Howard-Snyder (ed.), *The Evidential Argument from Evil*, 126–50. Bloomington: Indiana University Press.
- Wykstra, Stephen J. 2007. "CORNEA, Carnap, and Current Closure Befuddlement." *Faith and Philosophy* 24: 87–98.
- Wykstra, Stephen J. and Perrine, Timothy. 2012. "The Foundations of Skeptical Theism: CORNEA, CORE, and Conditional Probabilities." *Faith and Philosophy* 29: 375–99.
- Yandell, Keith. 1985. "The Problem of Evil and the Content of Morality." *International Journal for Philosophy of Religion* 17: 139–65.
- Yandell, Keith. 1989. "Gratuitous Evil and Divine Existence." *Religious Studies* 25: 15–30.