

JIM STONE

EVIDENTIAL ATHEISM

ABSTRACT. Here is a new version of the Evidential Problem of Evil. Theists claim that it is reasonable for atheists to believe that if God *did* exist, suffering would look just as it does now. I endorse this claim, however it cannot be deployed against my argument without the following epistemic principle: what we see makes *p* likely only if it is reasonable to believe it would be discernibly different if *p* were false. I demonstrate that this principle is mistaken. The paper also responds to objections from Alvin Plantinga and Peter Van Inwagen that God's existence is compatible with pointless natural evil. In particular, I argue that appeals to vagueness do not support the compatibility claim.

INTRODUCTION

The Logical Problem of Evil maintains that suffering logically precludes an omnipotent, omniscient, unsurpassably benevolent creator ('God,' henceforth); as suffering exists, God does not.¹ LPE has both global and local versions. The former claims that any suffering whatsoever precludes God; the latter that a particular sort of suffering (e.g. 'natural evil' – suffering due to earthquakes, cancer and other calamities that apparently flow from the operations of nature) entails atheism. No version is persuasive, however, for the claim that suffering logically precludes God's existence cannot be justified. God is compossible with any evil you choose if he must permit it in order to realize an outweighing good. Suppose for argument's sake that we cannot even *imagine* what that good might be – it would be hubris to conclude that it does not exist.

Partly for this reason,² interest in the theology of evil has recently devolved to an argument called the 'Evidential Problem of Evil,' which makes the weaker but still fiercely contested claim that the widespread and terrible suffering we see about us makes theism implausible. Each side is motivated by significant insights but fails to understand its piece of the truth well enough to deploy it deci-



Philosophical Studies **114**: 253–277, 2003.

© 2003 Kluwer Academic Publishers. Printed in the Netherlands.

sively. In what follows, I try to cut the Gordian knot by defending a simple new version of the Evidential Problem of Evil. Parts I and II set the stage for the new argument, which appears in III. Theists often claim that it is reasonable for atheists to believe that if God *did* exist, suffering would look just as it does now. This claim (which I will support) cannot be deployed against my argument without the following epistemic principle: what we see makes *p* likely only if it is reasonable to believe it would be discernibly different if *p* were false. Part IV demonstrates that this principle is mistaken. Parts V and VI answer objections that flow from the work of Alvin Plantinga and Peter van Inwagen, both of whom deny – for quite different reasons – the assumption that theism is inconsistent with pointless suffering.

In order to simplify my exposition, I will use the terms ‘probable’ and ‘plausible’ in a special way. My concern is the epistemic likelihood that a thesis has for a believer given her grounds for believing it. Let us say that *p* is ‘probable’ for *S* when *S* has grounds for believing *p* such that a fully rational person given them would, on that basis, affirm *p* – not merely judge that *p* is a better bet than not-*p*.³ A fully rational person would affirm *p* if and only if he, if asked whether *p* is true, would answer that it is or at least that he believes it is. As I mean the term, ‘grounds’ need not be (or involve) another belief. A fully rational person given a headache like mine would on that basis believe she had one. So *p* is ‘probable’ for *S* only if *S*’s grounds justify or warrant *S*’s believing *p*. Of course, *p* is ‘improbable’ for *S* if not-*p* is probable for *S*. The claim that *p* is probable *simpliciter* asserts that most everyone has grounds that make *p* probable for him.

Rational degrees of belief admit of a numerical representation (with larger numbers corresponding to stronger beliefs) that obeys the rules of the probability calculus (Skyrms, 1986). Where 1 is certainly true, 0 certainly false, and 0.5 no more believable than not, it seems plain that justified belief requires a high probability. I will use 0.9 in what follows; however my argument goes through if probability is set at 0.8. Note that ‘probable’ true belief is insufficient for knowledge – Gettier problems aside. Suppose I am one of a million ticket holders in a fair lottery. As I do not know who will win, I do not know who will lose; therefore I do not know

I will lose. Nonetheless the odds entitle me to believe that I will lose.

By contrast, I will say that *p* is 'plausible' for *S* when *S* has grounds for believing *p* such that a fully rational person given them would judge on that basis that *p* is a better bet than not-*p*. Probable beliefs must be plausible beliefs, but not vice versa. Finally, *p* is 'likely' for *S* if *p* is probable or plausible for *S*; and *p* is 'unlikely' for *S* if not-*p* is likely for *S*. This terminology will be helpful when I present the new version of the Evidential Problem of Evil.

I

Consider the thesis that every pleasant experience results in an outweighing pain. This is improbable given what we actually see. Of course some pleasures lead to such pains, but there are myriad enjoyments that by all appearances do not. Drinking cool water when one is thirsty, going for a pleasant walk, even lying down when one is weary, typically do not result in any pain at all, not to mention outweighing pains. These enjoyments often lead to more enjoyment, a sense of well-being, serenity, and so on. When we feel happy and serene we are usually less likely than otherwise to make other creatures suffer. In short, the world appears to contain billions of token enjoyments that do not result in any pains at all, not to mention outweighing pains. On balance they make the world a happier place.

Of course there may be intrinsic disutilities other than pain. Injustice and bad faith are disutilities even when they do not make the world less pleasant. But the claim that every pleasant experience results in an outweighing disutility is also improbable, because most of the myriad enjoyments I mentioned do not result in any disutilities as far as we can see. To the contrary, we see that they often lead to additional utilities. Suppose we divide disutilities into *positive* and *negative*: the former includes the actualization of positive intrinsic evils like pain and injustice, the latter the loss of positive intrinsic goods (e.g. the loss of pleasure or knowledge). Many people take pleasure in kindly acts; consequently their lives, which are sometimes extraordinarily productive of positive utilities, are also rich in apparently innocent pleasures. Such lives make

it quite improbable that every pleasant experience that does not result in an outweighing positive disutility results in an outweighing negative one.

Consider, then, thesis T:

Every pleasant experience results in an outweighing positive or negative disutility.

Let us say that experience *e* results in disutility *d* just in case both happen, and *d* would not have happened if *e* had not occurred;⁴ a disutility is *outweighing* when it is equal to or greater than the utility from which it results. T is disconfirmed by experience. When billions of instances of phenomenon X appear not to result in instances of Y, and there is no good reason to believe we would routinely miss the Ys that result from Xs if they were happening, then, barring new evidence, it is probable that many Xs do not result in Ys. (A 'good reason' is a likely hypothesis entailing that we would miss the Ys that result from Xs if they were happening.) In addition, if people have watched the human condition for millennia without finding candidate Ys that result from these Xs, that likelihood is increased.

Note that *d* can precede *e*. We must look for *d* 'upstream' as well as 'downstream' from *e* in evaluating T's likelihood. However the claim that, had I not watched the sunset, some bad thing that preceded my act would *not* have happened or some good thing that did not precede it *would* have occurred, is unlikely. One cannot find even a candidate positive preceding disutility that would not have happened had I not watched the sunset; such candidates are typically absent in the billions of sunset watchings that have occurred in the last few millennia. Also absent are candidate preceding positive utilities that *would* have been realized if the watchings had not happened. But what if the universe is deterministic? The immediate past would have been different had I not watched the sunset, and the past just before that, and so on indefinitely. Surely some preceding disutilities would have been lost! Suppose so. As there is just as much reason to believe that past utilities would have been lost too, we still have no candidate *outweighing* preceding disutility. But suppose that mostly (or only) disutilities would have been lost; probably plenty of sunset watchings are like that. Then probably plenty of sunset watchings are such that mostly (or only) utilities

would have been lost: all the more reason to conclude that T is false!

This account explains why we can be confident of the falsity of the thesis: “Every act of ear tugging (or hand clapping, or finger snapping) results in a significant disutility (or a significant utility, or a mystical experience of union with God).” The claims that (a) ear tuggings result in disutilities of which we have as yet no conception, (b) ear tuggings result in standard disutilities (e.g. hate crimes and cancer) in ways we do not discern, and (c) advancing science may one day persuade us that a and b are true, have the force of mere skeptical hypotheses unless they are motivated by a plausible theory (compare Alston, 1996). As there is no such theory, (a), (b), and (c) provide no reason to believe we would routinely miss the disutilities if they were happening or that we would fail to discern the connection between ear tuggings and standard disutilities if one did exist. Even if some of our actions result in standard disutilities in ways we do not discern, there is no reason to believe that ear tuggings do. Even if there are disutilities of which we have no conception, there is no reason to believe any of them result from ear tuggings. The claim that for all we know there are so many unknown disutilities that our sample is unrepresentative, so that we have no good reason to believe that ear tuggings (or apparently innocent enjoyments) do not result in some significant disutility or other, has the force of a mere skeptical hypothesis unless it is motivated by a plausible theory.

Such theses cannot be falsified conclusively, of course. We would have to conclusively verify that some ear tugging never results in a significant disutility, and no finite set of observations would suffice. Perhaps the disutility will happen in a billion years; maybe it has happened already in another galaxy far away. But disconfirmation need not be conclusive to make a thesis unlikely; consider “Some men are immortal.”⁵ To invert an example from Bertrand Russell, if I open a crate of oranges and the top layers appear to be in perfect condition, this provides a reason to think the remainder are in good condition, too. Similarly, it is more likely that the world will continue to look as it does concerning apparently innocent enjoyments than that it will become very nasty because of them in the distant future; and it is more likely that the rest of the universe is like the part of it we see than that things are nasty far away on

account of the apparently innocent pleasures we experience here. Indeed, we know of no laws or mechanisms that could link such pleasures to spatially or temporally distant disutilities. Hence what we see entitles us to conclude that some enjoyments appear not to result in outweighing disutilities – even though we cannot directly confirm the absence of the disutilities in all times and places.

Barring new evidence, surely it is wrongheaded to insist that any of these theses is as likely as its denial – T included. Indeed, it is reasonable to believe that many pleasant experiences do not result in outweighing disutilities; there is plenty of ‘pointless’ pleasure. Scientific accounts of reality proceed on the principle that things probably are as they appear to be throughout a multitude of observations by different observers over an extended period of time – unless we have a positive reason to believe otherwise (that is, a likely hypothesis which entails that things are *not* as they appear to be). This principle, ultimately, generates the Evidential Problem of Evil.

II

The mirror image of T is MT:

Every significant suffering that makes up the widespread and terrible suffering we see about us results in an outweighing positive or negative utility.

Let us say that *s*, a particular instance of suffering, *results* in utility *u* just in case both happen, and *u* would not have occurred if *s* had not happened;⁶ a utility is *outweighing* when it is equal to or greater than the suffering from which it results. (Of course, it is implicit that there is no disutility *d* that also results from *s* such that, when *d* is subtracted from *u*, *u* does not outweigh *s*.) *Pointless evil* is suffering that does not result in an outweighing utility. A *positive* utility involves the actualization of an intrinsic good like pleasure or justice; a *negative* utility involves the prevention of an intrinsic evil like injustice or pain. Hence MT entails that for every significant suffering *x*, there is an outweighing intrinsic good that would not have happened, or an outweighing intrinsic evil that would have happened, if *x* had not occurred.

The argument for MT's improbability is analogous to the argument for T's improbability. Here are three (actual) accounts of the sort of events that count cumulatively against MT. First, a twelve-year-old girl enters a hospital for cancer treatment, insisting that it is only for a short while; soon she will return to school and her friends. She maintains this steadfastly until the last day of her life, when she begins to weep inconsolably. Her parents' desperate efforts to comfort her are useless. She weeps inconsolably until she dies. Second, a 28-year-old woman, still struggling with the psychological ravages of being orphaned as a child, is haunted by the fear that something will prevent her from caring for her young daughter. Professional counseling finally persuades her to forget her fear. Then she suffers a stroke due to an aneurism, which leaves her brain-damaged, quadriplegic, and barely able to speak – though she understands exactly what has happened to her. She spends the rest of her life in a nursing home. Her daughter is devastated psychologically. Third, cheetahs usually suffocate their prey, but occasionally they disable an antelope and eat her alive.

Events like these, which appear not to result in any outweighing utility, positive or negative, happen daily by the thousands. Indeed, such sufferings often lead to more suffering – people go mad, children are brutalized, families go under.⁷ As human observers have been confronted by billions of such sufferings, and as there is no good reason to believe we would routinely miss the resulting outweighing utilities if they were happening or that they would be delayed until the distant future, then, barring new evidence, it is probable that many sufferings do not result in an outweighing utility. The claims that such sufferings result in utilities of which we have no conception, or that they are connected to standard positive or negative utilities in ways we do not discern, are mere skeptical hypotheses unless they are motivated by a plausible theory. The claim that there may be so many unknown utilities that our sample is unrepresentative, so that we have no good reason to believe the calamities I described above did not result in some outweighing utility or other, has the force of a mere skeptical hypothesis unless it is motivated by a plausible theory.⁸

Note that u may precede s. We must look for u 'upstream' as well as 'downstream' from s in evaluating MT's likelihood. However

the claim that, had the woman in the second account been spared the brain damage, some preceding intrinsic good would not have happened, is unlikely. Aneurisms are typically the result of non-genetic congenital defects. The closest possible world without that brain damage is one where the fetus developed normally. Nor is there any reason to think that an intrinsic evil that did not precede the stroke *would* have happened if the aneurism had not occurred; indeed, one cannot even find a candidate for what that evil might be. Consider also the daily toll of third-world children who die of infectious diseases. For each of them the closest world without that agony is one where he or she did not get infected. Where 'A' denotes the actual world, it is incredible that every last one of those worlds either lacks a preceding outweighing intrinsic good that A has or contains a preceding outweighing intrinsic evil A lacks. Finally, consider the improbability of the claim that in every case where cheetahs disabled an antelope and ate her alive, if instead they had killed her before eating her, then either some preceding intrinsic good would *not* have happened or some preceding intrinsic evil that did not happen *would* have happened. One cannot find even a candidate preceding good that would have failed to happen, or a candidate preceding evil that did not happen but would have, if the cheetahs had killed the antelope before they ate her.

In short, barring new evidence, MT is improbable for the same kind of reasons that T is improbable. MT is another improbable claim about the course of nature of the sort we considered in section I, above.

III

Here is the new version of the Evidential Problem of Evil (call this new argument 'EPE'):

1. What we see of suffering makes MT improbable.
2. We are justified in believing that theism entails MT.

The force of 2 is that it is probable for us that theism entails MT. I take a sentence of the form 'p entails q' to express the thought that there is no possible world where p and not-q. My grounds for believing that p entails q might include a priori intuitions, the

testimony of experts, and arguments (e.g. “An omniscient being would know of pointless suffering, an omnipotent one could eliminate it, an unsurpassably benevolent one would wish to do so; hence God would eliminate pointless suffering”). Also my grounds might include observations and experiments (for instance, these support the claim that “There is water” entails “There is H₂O”). My grounds make it probable for me that *p* entails *q* when a fully rational person given them would on that basis affirm that *p* entails *q*. (2 will be reconsidered in part V, below.)

3. For any propositions *p* and *q*, if what we see makes *q* improbable, and we have the justified belief that *p* entails *q*, then what we see makes *p* unlikely.

Equivalently: (3*) if what we see makes *t* probable and we have the justified belief that *t* entails *u*, then what we see makes *u* likely.⁹ But what if *u* is already likely on other grounds? The consequent of 3* should read: ‘then what we see makes *u* likely – unless *u* is already likely on other grounds.’ And the consequent of 3 should read: ‘then what we see makes *p* unlikely – unless *p* is already unlikely on other grounds.’ I will err on the side of theism and assume for argument’s sake that it is *not* already unlikely on other grounds; so the ‘unless’ clause will be omitted in what follows.¹⁰

3 entails (by substitution)

4. If what we see of suffering makes MT improbable, and we have the justified belief that theism entails MT, then what we see of suffering makes theism unlikely.

1, 2, and 4 entail

5. What we see of suffering makes theism unlikely.

As human observers have been confronted by billions of instances of apparently pointless suffering, and things probably are as they appear, probably there is pointless suffering – unless a likely hypothesis entails that things are *not* as they appear. Probably theism logically precludes pointless suffering. It follows that what we see of suffering makes theism unlikely. More obscurely: As ‘not-MT’ and ‘not-MT entails atheism’ each have a probability of 0.9, their conjunction has a probability of 0.81; as that conjunction entails atheism, and likelihood is preserved under entailment according to the probability calculus, theism has a likelihood of no more

than 0.19 – a long shot. Suppose, however, I realize that Anselm’s Proof works. As I have a good reason to believe theism, I have a good reason to believe that we routinely fail to see the outweighing utilities that in fact always do result from apparently pointless sufferings; therefore I have a good reason to deny 1. In short, EPE is defeasible. Of course a good reason to believe theism, if it makes theism only slightly plausible, might not overturn 1.¹¹ We can safely say that the force of 5 is this: what we see of suffering makes theism unlikely – unless we have a good reason to believe that God exists.

IV

Line 3 of EPE might seem uncontroversial, yet here I diverge significantly from the contemporary discussion of the problem of evil. Stephen J. Wykstra maintains that we are entitled to claim ‘it appears that p’ only if the following Condition of Reasonable Epistemic Access (CORNEA) is satisfied:

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 by her. (Wykstra, 1990, p. 152)

Put simply, the force of CORNEA is that what we see or experience makes *p* likely (hence gives us ‘epistemic access’ to how things are) only if it is reasonable to believe that it would be discernibly different if *p* were false. But, Wykstra claims, it is reasonable to believe that if God exists the apparently pointless suffering we see about us would *still* look pointless. For God can conceive of goods we cannot imagine and these would be the outweighing goods the suffering secures. Given CORNEA, therefore, what we see of suffering cannot support atheism – for it is reasonable to believe that if theism were true, things would look just as they do.

William Rowe, the principal champion of the Evidential Problem of Evil, agrees that God would certainly conceive of goods ‘beyond our ken.’ Nonetheless the hypothesis of theism gives us ‘no reason whatever to suppose *either* that the greater goods in virtue of which he permits most sufferings are goods that come into existence far

in the future of the sufferings we are aware of, *or* that once they do obtain we continue to be ignorant of them and their relation to the sufferings' (Rowe, 1990, p. 165). Rowe concludes that it is reasonable to believe that we would probably not see vast quantities of apparently pointless human and animal suffering if God existed.

Let me offer en passant an argument on Wykstra's behalf. We might well address the question "How would suffering look to us if God existed?" by asking another one: "Is the closest God-occupied possible world one where suffering looks different (either because there is less of it or because more of it produces recognizable goods), or one where it looks the same and is necessary for God's securing outweighing utilities (either goods beyond our ken, future goods, or standard utilities such that the connection between the suffering and the utilities is yet to be discerned)?" On the face of things the second world is closer because it more closely resembles the actual world. Otherwise the issue is undecidable. We can hardly estimate what an omniscient creator's good-securing strategies would be. Either way it is *not* reasonable to believe that suffering would look different if God exists. This insight often underlies theistic objections to the Evidential Problem of Evil.

What matters most for our purposes, however, is that Rowe accepts CORNEA. Consider the thesis that Rowe calls 'Expanded Standard Theism': God exists and the goods in question will be realized only at the end of the world. 'Restricted Standard Theism' is simply the thesis that God exists. Rowe writes: "EST is not rendered unlikely by the items that render RST unlikely. And this, for precisely the reason that Wykstra so clearly and carefully sets forth" (Rowe, 1990, p. 166). This is a mistake. As EST plainly entails RST, given 3 it follows that EST is unlikely if RST is improbable – unless we have a good reason to believe EST. The items that render RST unlikely render EST unlikely, too. In short, CORNEA is false if 3 is true.

Of course Wykstra and Rowe might well respond that, as CORNEA is true, 3 is false. CORNEA leads to absurdity, however. Consider the hypothesis that the Earth is flat (call it 'F'). Consider just one piece of evidence against F: the Earth looks spherical from space. This by itself makes not-F probable – unless there's a likely

hypothesis which entails that the Earth is not as it appears from space. But now consider thesis F^* :

The Earth is flat, and (a) a force field around it bends light so that the Earth looks spherical from space.

Given CORNEA, the fact that the Earth looks spherical from space isn't evidence for not- F^* . For if F^* were true, the Earth would look spherical from space. On the other hand, CORNEA explains why not- F is probable given what we see. As there is no good reason to believe (a), it is reasonable to believe that, if F were true, the Earth would *not* look spherical from space. But it is paradoxical to claim that, while what we see makes F improbable, it does not make at all unlikely the conjunction of F and (a). F^* can be no more likely than F , plainly.

Friends of CORNEA might agree that F^* is unlikely, but insist that this flows from a source other than our evidence for not- F . After all, F^* is the conjunction of two bold claims; such a conglomeration is implausible on its face. Indeed, as what we see would be just the same if F^* were true, *how* can it make F^* unlikely?¹² Here is how. F^* is a conjunction: ($F \ \& \ (a)$). Hence its negation (not- F^*) is a disjunction: (not- $F \vee$ not-(a)). As what we see makes not- F likely, it makes this disjunction at least as likely as not- F ; therefore it makes F^* at least as unlikely as F . Consequently not- F^* is rendered likely by the items that render not- F likely – these are evidence for not- F^* too – even though it is unreasonable to believe they would be discernibly different if F^* were true. It follows that CORNEA is mistaken.

Here is the argument set out more formally:

- d. That the Earth looks spherical from space makes not- F likely, and CORNEA helps to explain this fact.
- e. Necessarily, for any p , observations that make p likely make any disjunction that includes p as a disjunct at least as likely as p .

d and e entail

- f. The observations that make not- F likely make it likely that (not- $F \vee$ not-(a)) – which is plainly equivalent to not- F^* .

f entails

- g. Those observations make not-F* likely.
- h. If CORNEA is true, g is false (for plainly the Earth would look spherical from space if F* were true).

g and h entail

- i. CORNEA is false.¹³

The upshot: Suppose we know that if God existed, suffering would appear just as it does now. This no more warrants denying that what we see of suffering makes theism unlikely than knowing that the Earth would look spherical from space if F* were true warrants denying that what we see of the Earth makes F* unlikely.¹⁴ This is the insight that underlies EPE. Now we can diagnose the impasse in the theological debate. Both sides accept that (i) MT's improbability would make theism implausible, and (ii) CORNEA is correct. (i) and (ii) conjointly entail that what we see of suffering makes MT improbable only if it is reasonable to believe that what we see would be different if God exists. As atheists are convinced (correctly) that MT is improbable given what we see, they conclude that it is reasonable to believe things would look different if God exists. As theists are convinced (correctly) that what we see would likely look the same if God exists, they conclude that it does not make MT improbable. Hence much fruitless wrangling that conflates the question "Does what we see make MT improbable?" with "Would things look different if God exists?" CORNEA (or, rather, the intuition it expresses) is the culprit, I maintain. Once we see that the fact that things would look the same if God exists does *not* count against MT's improbability, which still makes theism unlikely, the way forward is clear.

V

Let me explain why line 2 of EPE is false as it stands. Earlier I distinguished global from local versions of LPE. The first premise of any such argument cannot be justified, I argued. A more interesting response to global LPE is the 'Free Will Defense' (FWD): God must allow moral evil if he allows us moral free will – as free will is a

great good, global LPE's premise is false. I will err on the side of theism and suppose that FWD succeeds.

Note that moral evil is pointless. Suppose I cause suffering *s* by performing a wicked act at *t*. According to FWD, the prime candidate for a resulting outweighing good is my exercising at *t* my moral free will. But the closest world where *s* does not happen is one where I freely refrain from the wicked act; so I exercise my free will in that world too. As *s* does not result in an outweighing utility, the world would simply have been better if *s* had not happened. Nonetheless God is justified in allowing *s*, for he deprives me of my freedom in the more distant world where he interferes with me; hence that world is less good than the actual world. As theism is consistent with pointless suffering, it evidently does *not* entail MT.

We must distinguish moral from natural evil. Moral evil is the evil men do; natural evil is the suffering that apparently flows from the operations of nature (roughly, natural evil is the rest of suffering). Let me stipulate that MT concerns only natural evil. As my argument for MT's improbability adverted only to natural evil, its force is undiminished. Does theism entail MT so construed? Not if we accept FWD. Alvin Plantinga argues that FWD applies to the local version of LPE that appeals to natural evil (Plantinga, 1974). In possible world *w*, God gives free will to non-human persons, too; they deliberately cause cyclones, strokes, and cancer. In *w* natural evil is the evil they do (*w* is a 'demon world,' we might say). As their free will is a great good, God must allow natural evil as its price; therefore God and natural evil are compossible. Note, however, that the suffering the demons cause in *w* is pointless. They would have been just as free if they had behaved themselves, so it results in no outweighing utility. As theism evidently does *not* entail MT (construed to concern only natural evil), line 2 is still false.

My response is to modify EPE. A 'zany' world involves, relative to our probable beliefs about the actual world, a bizarre scenario, an improbable ontology, or a plainly factually mistaken claim. Demon worlds are zany, for instance. So as not to beg the question against theism, we will proceed on the (reasonable) supposition that a world's having God in it does not make it zany. *C* is the conjunction "God exists and we are not in a zany world."

1. What we see of suffering makes MT improbable.
2. We are justified in believing that C entails MT.

Earlier I wrote:

God is compossible with any evil we choose if he must permit it in order to realize an outweighing good. Suppose for argument's sake that we cannot even *imagine* what that good might be – it would be hubris to conclude that it does not exist.

Is it not also hubris to infer from the fact that we cannot imagine one, that probably there is no God-occupied non-zany world containing pointless natural evil? I do not think so. Such a world would have to contain an extraordinary good that does *not* result from the suffering, which God would nonetheless destroy if he interfered with the suffering, a good that does not involve the free will of demons or any other zany scenario. It is reasonable to believe that C entails MT, I submit – at least until we are told what that good might be.

3. For any propositions p and q, if what we see makes q improbable, and we have the justified belief that p entails q, then what we see makes p unlikely.

(That is, if 'not-q' and 'not-q entails not-p' each have a likelihood of 0.9, then their conjunction has a likelihood of 0.81 and p has a likelihood of no more than 0.19.)

3 entails (by substitution)

4. If what we see of suffering makes MT improbable, and we have the justified belief that C entails MT, then what we see of suffering makes C unlikely.

1, 2, and 4 entail

5. What we see of suffering makes C unlikely.

(That is, C has a likelihood of no more than 0.19.)

6. If one of the two conjuncts in a conjunction with a likelihood of 0.19 or less is probable, then the other conjunct has a probability of no more than 0.2111 . . .
7. Probably we are not in a zany world.

5, 6, and 7 entail

8. Theism is a long shot (0.2111 . . .) given what we see of suffering.

Probably we are not in a zany world, so a good reason to believe C must be a good reason to believe theism.

VI

Peter van Inwagen would deny that it is reasonable to believe that C entails MT. He writes:

It is not very plausible to suppose that there is a way in which evil could be distributed such that (i) that distribution of evil would serve God's purposes as well as any distribution could and (ii) God's purposes would be less well served by *any* distribution involving less evil. (One might as well suppose that if God's purposes require an impressively tall prophet to appear at a certain place and time, there is a minimum height such a prophet could have.) (van Inwagen, 1988, p. 199)

Here is why I find this ingenious objection unpersuasive.

1. Consider the statement "Isaiah, who is 7-feet-tall, is impressively tall; and for any height an impressively tall prophet has at a certain time and place, it can be reduced without his ceasing to be impressively tall." This generates a classical sorites paradox, for it apparently entails that if we shrink Isaiah by tiny increments he will be impressively tall when he is 3-feet-high. The Stoics responded by denying the minor premiss of such arguments: there is indeed a minimum height such a prophet can have, even if we cannot know it. The view that sorites paradoxes pose an epistemic problem has been revived recently by Roy Sorensen and Timothy Williamson. Of course a sharp boundary between an impressively and a merely tall prophet seems incredible. As Williamson writes: "No logic for a vague language seems wholly congenial to common sense. . . . The truth about vagueness must be strange" (Williamson, 1994, p. 166). In addition, the Stoic's solution preserves the principle of bivalence, thereby preserving classical truth-conditional semantics and classical logic – which other accounts of vagueness typically compromise.

Arguably the 'epistemic view' is the worst there is, except for all the others. Consequently the claim that there is a minimum height

the impressively tall prophet appearing at that time and place could have is less silly than it first appears. Indeed, as the Stoic's view has become a leading view, the claim that there is no minimum quantity of suffering required to serve God's purposes is harder to motivate by an appeal to vagueness.¹⁵

2. There is a relevant disanalogy between magnitudes of suffering and magnitudes of time and space. Consider duration. A billionth of a second is a duration no less; indeed, it is itself infinitely divisible. However a billionth of a second of suffering is not suffering at all. There is an essential subjectivity to suffering: as there is nothing it is like to suffer for a billionth of a second, there can be no span of suffering so brief. Nor is there anything it is like to suffer for a billionth of a second less. Divisions in the duration during which suffering occurs outrun divisions in the temporal magnitude of the suffering itself. Shifting from duration to intensity, consider Oedipus's suffering as he staggers off, a blind beggar, on his way to Colonus. It makes no sense to talk of a billionth of a part of the intensity of that suffering, or of reducing its intensity by a billionth of a part. In short, along the dimensions where there is more or less of it, suffering appears *not* to be infinitely divisible. It comes in gobs, not microns, and differences in its magnitude are perceptible in principle.

This disanalogy further weakens van Inwagen's objection. Certainly it is hard to believe that reducing the impressively tall prophet's height by a billionth of an inch would alter his power to impress us. But if *per impossibile* height did come only in perceptible units, there might well be no part of his height so small that removing it would leave his power to impress unaltered. Suppose, then, that God allows natural evil because it results in an outweighing good. S is the sum of suffering in the universe. As suffering comes in perceptible gobs, not microns, the prophet analogy provides no reason to believe that some proper part of S would have served God's purposes just as well. Indeed, given the 'gobyness' of suffering the claim is plausible that no part of S is so small that removing it would leave unaltered the powers of the remainder; otherwise God would have removed it.

The ‘gobbiness’ of suffering also blocks the non-sorites argument that, supposing a certain quantity q of suffering is *almost* enough to realize the Good (so that any quantity greater than q will suffice), for any quantity greater than q there is an infinity of lesser sufferings between it and q that will do as well.¹⁶

3. Consider this example: God allows suffering so as to produce a certain set of utilities U at time t . U consists of a number of people, including Jo, Ted, Flo, . . . each accepting Jesus whole-heartedly as his or her saviour. Consider a spectrum of close possible worlds such that the suffering of our world is reduced slightly from one to the next. U is realized at t in the actual world; no member of U happens at t in the last. Throughout the range there is always a fact of the matter as to whether every member of U happens at t . As suffering is required to bring U about at t , somewhere in that range of worlds there is one where a member of U does *not* happen at t ; so a minimum amount of suffering is required to achieve God’s purpose.

As concerns about vagueness motivate van Inwagen’s objection, we ought to be able to make it more plausible by introducing vagueness into this example. Suppose that God’s purpose is that U be realized sooner or later, or that he is willing to settle for each of us more or less accepting Jesus. Throughout the range of cases where suffering is reduced and either U takes longer to be realized or is realized to a diminished degree, surely God’s purposes are being accomplished *less well* – unless we insist that he does not even prefer that we accept Jesus whole-heartedly and sooner rather than half-heartedly and later. Vague purposes can be satisfied more or less. As what God prefers is better than what he is merely willing to settle for, as we move along the range of cases where suffering is diminished and God’s preferences are less well satisfied, in each case a utility is lost. Were that good not worth the suffering from which it results, an unsurpassably benevolent being *would* prefer its loss. As each suffering results in an outweighing utility, it has a point after all.

4. Let us try again. As the ‘epistemic view’ is uncongenial to Van Inwagen’s objection, set it aside. In borderline cases there is no fact

of the matter as to whether Isaiah is bald, say; hence the claim that he is bald is neither true nor false (we can say it is 'indefinite'). Now my original description of the example was tendentious because I assumed there is always a fact of the matter as to whether U is accomplished at t. Suppose instead that there is a range of cases where it is definite that I have accepted Jesus whole-heartedly, and, as suffering diminishes, another range where it is indefinite whether this is so. Note, however, that God's purpose is served less well in the latter cases – unless we insist that he does not prefer that it is true rather than not true that I accept Jesus whole-heartedly. If God wants me to accept Jesus whole-heartedly, the fact of the matter matters.¹⁷

As we want to deploy vagueness to support Van Inwagen's objection, let us say that the border between definite and indefinite cases is vague. Otherwise a minimum quantity of suffering is required for God's purpose. Let 'P' name the proposition that I have accepted Jesus whole-heartedly. There is a range of cases where it is indefinite whether P has a truth value. But consider. In such cases it is neither true nor false that P has a truth value. It follows that P is not true. For suppose that P is true – it follows that it is false that it is neither true nor false that P has a truth value. Therefore God's purpose is served less well. Again we must say we have reached a vague border. There is a range of cases where it is indefinite whether it is indefinite whether it is definite that I have accepted Jesus whole-heartedly. Again it follows that P is not true – an inference that iterates ad infinitum, if necessary. If there is an infinite series of higher-order blurry borders, at every one of them God's purposes are served less well.

5. Finally, as suffering is not infinitely divisible, there cannot be an infinite series of higher-order blurry borders. We began with a range of cases where it is definite that I accept Jesus whole-heartedly; as the quantity of suffering in the world was reduced we moved into a range of indefinite cases. Call the entire range of cases R. We can represent the reductions in suffering as a series of contiguous marks, each having an equal indivisible extension, which together constitute a line. As we move from left to right along the line there is less suffering; as each case in R depends upon the quantity of suffering in the world, each mark corresponds to just one case in R.¹⁸ Only

a small part of the line corresponds to cases in the vague border between definite and indefinite cases. As we ascend to higher-order borders a diminishing area of the line corresponds to the cases in them; otherwise there will be no cases in R where it is definite that I accept Jesus whole-heartedly. As suffering is not infinitely divisible, we must sooner or later reach a part of the line that includes only one indivisible mark. By hypothesis this corresponds to a border between definite and (some range of) indefinite cases. As the border is crossed if we remove this bit of suffering, a minimum quantity of suffering is required for God's purpose.

Given these difficulties, I submit that van Inwagen's objection is without force unless it can be motivated by specific examples of God's purposes that would be satisfied just as well if there were less suffering. (Of course, implicit in this request is that the purposes aren't counter-intuitive or puzzling ones for God to have.) These are hard to find. van Inwagen writes (concerning the evil era brought about by man's fall): "If there is a purpose that is served by allowing 'the age of evil' to have a certain duration, doubtless the same purpose would be served if the age of evil were cut short by a day or a year or even a century" (199). However no clear suggestion is offered as to what God's purpose might be.

Might retribution serve? Suppose God punishes certain malefactors by giving them terminal cancer. No quantity of suffering is perfectly calibrated to what the malefactor deserves; God's concern is only that she suffer mightily. A difficulty is how to square unsurpassable benevolence with giving people cancer out of retribution. God would deliberately inflict suffering only for a beneficent end, I submit – to cause the malefactor to repent and reconcile herself to him, say – but then he knows how much it will take and when to quit.

Let us accept the suggestion anyway. God has at least one purpose that would be satisfied just as well if there were less suffering. As plenty of apparently pointless natural evil befalls non-malefactors, we need only restrict MT to that sort of suffering (e.g. suffering befalling good people, bad people who have already suffered terribly, small children, and non-human animals) and EPE is back in business. The real challenge, therefore, is to specify purposes of an unsurpassably benevolent creator that would be satis-

fied just as well if there were less suffering – for every sort of natural evil of which there appear to be plenty of pointless instances. Until then the objection to 2 is unmotivated.

To conclude: Unless we have a good reason to believe that something exists that not only is all powerful and all knowing but is also the most benevolent being possible, the widespread and terrible suffering we see about us makes theism a long shot.¹⁹

NOTES

¹ Why not evil of any sort, including moral evil producing no suffering (e.g. I break my promise to grandmother to put flowers on her grave)? Painless evil does not make anybody shake her fist at the sky. This paper's focus is the evil that makes people atheists. 'Theism' will denote the view that God exists.

² A more interesting reason that LPE has been rejected ('The Free Will Defense') will be presented in part IV, below.

³ The notion of a 'fully rational person' is borrowed from Paul Draper's tentative account of epistemic probability; see Draper, 1989. I do not believe that anyone has yet given an adequate account of epistemic probability, only that this one works well enough to serve this paper's purpose. I suppose that the fully rational person is inclined to answer our questions sincerely.

⁴ A *possible world* is a way the universe might have been. A possible world W is 'closer' to world X than is world Y if W resembles X more than Y does. The world where I am a politician is closer to the actual world than the one where humans never evolved. I will follow the Lewis-Stalnaker account of counterfactuals, according to which 'Disutility d wouldn't have happened if pleasant experience e hadn't happened' means 'd doesn't happen in the closest possible world where e doesn't happen.' Hence e *results* in d just in case e and d both happen and d fails to occur in the closest possible world where e does not happen. As d may precede e, 'results' cannot be defined in terms of causation.

⁵ Even if every human being has died except for you and me, and we are about to commit hari kari, maybe new men will emerge one day and one of them will be immortal. "Some men are immortal" cannot be conclusively falsified (or conclusively verified); however we have considerable inductive evidence that it is false.

⁶ As explained in note 4, above, I will follow the Lewis-Stalnaker account of counterfactuals. Consequently s *results* in u just in case s and u both happen and u does not occur in the closest possible world where s does not happen. Therefore the relation between s and u is a weak one. Michelangelo's mother's labor pains result in my delight in seeing the Sistine Chapel. In addition, while her giving birth is causally implicated (as a *sine qua non*) in my pleasure, the accompanying pains are a side effect. This has two consequences. First, if God insures that every suffering is robustly causally involved in the production of an outweighing good, not a side effect or merely a causal *sine qua non*, MT follows a fortiori. The

reverse entailment does not obtain, however. Hence the claim “We’re justified in believing that theism entails MT” is rendered more plausible. Insuring MT’s truth is the least one would expect of an unsurpassably benevolent being. Second, MT is itself more plausible. On the face of things, there is less pointless suffering than there would be if side effects and mere necessary conditions did not ‘result’ in downstream utilities absent in the closest possible worlds without them. In defining ‘results’ as I do, I (a) minimize the demands on the Almighty, and (b) err on the side of theism.

⁷ The response that such calamities provide those directly affected (including loved ones) the opportunity to develop character or come to terms with the human condition leads to multiple difficulties. (1) The positive disutility often substantially outweighs the alleged utility, as in the first two accounts. (2) Often those directly affected have already come to terms with the human condition. Many on our planet are profoundly accepting of suffering, having suffered all their lives. (3) Many are too mentally ill, senile, or moribund to develop their characters. Others are overwhelmed and go mad. Often there are no loved ones. (4) Many die in agony before they can come to terms with anything, as when orphans die while still toddlers, babies starve to death, or a village is swallowed by an earthquake. (5) Calamities often reduce drastically the opportunities for moral growth; e.g. the opportunity for character development provided by being orphaned is often quite poor. (6) The response does not apply to the third example.

⁸ As philosophy and the sciences have not discovered new intrinsic goods – with the 2,500-year-old exception of the love of wisdom – the claim that they may well acquaint us with new ones is unwarranted (compare Alston 1996).

⁹ Epistemic closure principles (of the general kind ‘if *t* has a certain epistemic status and *t* bears some entailment relation to *u*, then some epistemic status can be inferred for *u*’) are controversial in recent epistemology. The strongest attacks have concerned closure for knowledge (‘*S*’s knowing *p* and knowing that *p* entails *q* is sufficient warrant for *S* to know *q*’), not likelihood, and typically depend upon accounts of knowledge which are arguably less intuitive than the principles themselves. It is hard to believe that we will ultimately abandon the principle that knowing an argument is valid and its premisses true is sufficient warrant to know the conclusion; and it is strongly counter-intuitive to insist that premisses ‘*t*’ and ‘*t* entails *u*’ are probable for me but my consequent affirmation of *u* is wholly without warrant. For more on closure principles see (Stone, 2000).

¹⁰ Arguably theism is a bad bet even before we consider evidence, for it is equivalent to a conjunction: “Something is omniscient, and something is omnipotent, and something is unsurpassably benevolent, and something created the Universe, and these somethings are one and the same.” A consequence of the probability calculus is that where *p* and *q* are independent, $\Pr(p \ \& \ q) = \Pr(p) \times \Pr(q)$. Some of theism’s conjuncts are *not* independent (e.g. the first conjunct makes the second more likely (though the third is independent of the others)). But even if we assign the first conjunct a probability of 0.9 and suppose that the conditional probability of each of the remaining conjuncts is also 0.9, theism’s initial probability is 0.45. As my concern is how what we see of suffering affects

theism's epistemic likelihood, I will set this aside. (I do assume, however, that theism's antecedent probability is not so great as to by itself make theism likely.) My point is this: If we suppose for argument's sake that theism is not a bad bet to begin with, what we see of suffering makes it one. Consequently if theism is a bad bet to begin with, what we see of suffering makes it worse.

¹¹ Suppose my good reason to believe theism gives it a probability of 0.52. Suppose the claim that theism entails MT has a probability of 0.9. MT has a probability of only 0.468, so it remains unlikely. Consequently I lack a good reason to believe that we routinely miss the outweighing utilities that always flow from apparently pointless sufferings.

¹² I am indebted to Jeffrey Brower for this entire objection.

¹³ A better principle might include the condition: "What we see increases s's likelihood if it is reasonable to believe that it would be discernibly different if s were false." CORNEA goes wrong because it substitutes 'only if' for 'if.' Defending and refining such principles is beyond this paper's scope.

¹⁴ Consider conjunction C: "Theism & MT." As MT is improbable given what we see, so is C – even though it is reasonable to believe that things would look as they do if C were true. Further, if theism entails MT it entails C; theism and C are logically equivalent. So if 2 is probable, theism is unlikely.

¹⁵ A note on terminology. I use the term 'border' to denote any boundary between one range of cases and another, whether it is sharp or vague. A vague (or 'blurry') border is one that involves borderline cases. 'Borderline' cases are ones where we do not know whether Isaiah is bald, say (though suppose we know how many hairs he has on his head). Certainly there are borderline cases on the Stoic's view, but there is always a fact of the matter as to whether Isaiah is bald. If we were not ignorant of such facts there would be no borderline cases. On the 'semantic view,' to which we will shift presently, borderline cases exist because in them there is no fact of the matter for us to know; hence the proposition that Isaiah is bald has no truth value.

¹⁶ Thanks to Chase Wren for pointing this out.

¹⁷ Generally if someone judges it desirable that p, she prefers that p is true rather than not true, even if p is not false. Indeed, one proposition that I prefer be true rather than not true is that I get what I desire. This presumption also goes for God. For simplicity's sake I omitted from my example vagueness concerning the time at which I accept Jesus whole-heartedly. Suppose there is a range of cases where it is definite that I accept Jesus at t, followed by a range of cases where it is neither true nor false that I accept Jesus at t, followed by a range of cases where it is true that I accept Jesus at a time later than t. As God (by hypothesis) judges it desirable that I accept Jesus at t, his purpose is less well served in the latter two ranges of cases. Finally, the argument of this section does not depend on God's wanting me to accept Jesus *whole-heartedly* – it works just as well if God simply wants me to accept Jesus.

¹⁸ We can guarantee isomorphism by the following strategy. We are considering a range of close possible worlds, starting with the actual world, such that suffering is diminished by its smallest unit from one to the next. Call that range S. There is

also a range of cases where it is true that I accept Jesus whole-heartedly; as the quantity of suffering in the world is reduced we move into a range of indefinite cases. Call this entire range of cases R. Each case in R is *identified* with one of the worlds in S. So the actual world is a case where I accept Jesus whole-heartedly, and the closest world with slightly less suffering, where presumably I also accept Jesus whole-heartedly, is a different case of my accepting Jesus whole-heartedly. As cases of accepting Jesus are individuated by worlds in which suffering is always slightly less than in the preceding world, each case in R corresponds to just one mark on the line and vice versa.

¹⁹ Thanks to Jeffrey Brower, Skip Larkin, Thad Metz, Philip L. Quinn, Kent Staley, Eleonore Stump, and Chase Wren for comments. Special thanks to Judith Crane.

REFERENCES

- Adams, M.M. and Adams, R.M. (eds.) (1990): *The Problem of Evil*, Oxford University Press.
- Alston, W.P. (1996): 'The Inductive Argument from Evil and the Human Cognitive Condition', in Howard-Snyder (1996) (pp. 97–125).
- Draper, P. (1989): 'Pain and Pleasure: an Evidential Problem for Theists', *NOUS* 23, 12–29.
- Howard-Snyder, D. (1996): *The Evidential Argument From Evil*, Indiana University Press.
- Howard-Snyder, D. (1996): 'The Argument From Inscrutable Evil', in Howard-Snyder (1996) (pp. 286–310).
- Howard-Snyder, D. and Howard-Snyder, F. (1999): 'Is Theism Compatible With Gratuitous Evil?', *American Philosophical Quarterly* 36, 115–129.
- Plantinga, A. (1994): *The Nature of Necessity*, Ch. 9 (pp. 164–193), Clarendon Press.
- Rowe, W.L. (1979): 'The Problem of Evil and Some Varieties of Atheism', *American Philosophical Quarterly* 16, 335–341. Reprinted in Adams (1990): 126–137.
- Rowe, W.L. (1984): 'Evil and the Theistic Hypothesis: A Response to Wykstra', *International Journal for Philosophy of Religion* 16, 95–100. Reprinted in Adams (1990): 160–167.
- Skyrms, B. (1986): *Choice and Chance: An Introduction to Inductive Logic*, 3rd edn., Wadsworth: Belmont.
- Sorensen, R.A. (1988): *Blindspots*, Oxford: Clarendon.
- Stone, J. (2002): 'Skepticism as a Theory of Knowledge', *Philosophy and Phenomenological Research* LX(3), 527–546.
- Stump, E. and Murray, M.J. (eds.) (1999): *Philosophy of Religion: The Big Questions*, Blackwell Publishers.

- Van Inwagen, P. (1988): 'The Magnitude, Duration, and Distribution of Evil: A Theodicy', *Philosophical Topics* 16(2), 161–187; also in Stump and Murray (1999): 195–210. Page references to Stump and Murray.
- Van Inwagen, P. (1991): 'The Problems of Evil, Air, and Silence', in J.E. Tomberlin (ed.), *Philosophical Perspectives* 5, *Philosophy of Religion*, Atascadero, CA: Ridgeview Publishing Co.; also in Howard-Snyder (1996): 151–174.
- Williamson, T. (1994): *Vagueness*, Routledge: London and New York.
- Wykstra, S.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. Reprinted in Adams (1990): 138–160.

University of New Orleans

Copyright of Philosophical Studies is the property of Kluwer Academic Publishing / Academic and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.