(Draft from 01/18/17)

Skeptical Theism and The Paradox of Evil*

Luis R.G. Oliveira

Abstract: Given plausible assumptions about the nature of evidence and undercutting defeat, it has seemed to many that the force of the *evidential problem of evil* depends on *skeptical theism* being false. I think this dialectic is mistaken. In this paper, I argue that there is a way of understanding the evidential problem of evil where it is compatible with skeptical theism. I suggest a way of defending William Rowe's famous 1979 argument that makes it depend on the evidential support provided by the collection of instances of apparently pointless suffering in a way that is compatible with each particular instance failing to provide any support at all. I call this result *the paradox of evil*.

Introduction

According to the *evidential problem of evil*, our seeing no justifying-reason for many instances of suffering is sufficient evidence for the belief that the traditional God does not exist. According to *skeptical theism*, however, it is not at all likely that we would see a justifying-reason for instances of suffering, were such a God and such reasons to really exist. Given plausible assumptions about the nature of evidence and undercutting defeat, it has seemed to many that the force of the evidential problem of evil therefore depends on skeptical theism being false. If we cannot expect to see God's justifying-reasons, were Him and them to truly be there, then our not seeing them can hardly count as evidence against His existence.

In this paper, I argue that there is a way of understanding the evidential problem of evil where it is in fact compatible with skeptical theism. I show that skeptical theism blocks the evidential problem of evil only given certain natural assumptions about how the evidence from evil accrues; I show that these assumptions are not essential to the problem; and I show which alternative assumptions can take its place. My aim here, however, is not that of *endorsing* the evidential problem of evil on the basis of these alternative assumptions. My aim is simply that of showing

*I'm grateful to Steve Wykstra, Timothy Perrine, Joe Milburn, Brendan de Kenessey, Scott Hill, Kristian Olsen, Bob Gruber, and Heidi Furey for helpful comments and discussion of previous drafts.

how the stalemate between the many who defend skeptical theism and the many who criticize it can be altogether sidestepped.

Here is how I proceed. I begin, in section 1, by clarifying two essential features of William Rowe's justly famous original formulation of the evidential problem of evil. Next, in section 2, I articulate what I call its *inductive justification*, which I argue is widely presupposed by Rowe commentators, according to which Rowe's argument depends on the accumulation of little bits of evidential support from particular instances of apparently pointless suffering. Then, in section 3, I argue that skeptical theism, properly formulated, resists Rowe's argument by denying that these particular instances provide even a modicum of support against God. With this dialectic clarified in the background, in section 4, I turn to a discussion of the preface paradox in order to identify a paradoxical evidential feature of certain inductive bases that is often ignored. Based on this analysis, in section 5, I suggest an alternative justification for Rowe's original argument. On what I call its *collective justification*, Rowe's argument turns on the evidential support provided by the collection of instances of apparently pointless suffering in a way that is compatible with each particular instance failing to provide any support at all. Drawing on the evidential dimension of the preface paradox, I call this result *the paradox of evil*. I conclude, in section 6, by arguing that skeptical theism, based as it is on a claim about our cognitive limitations, is compatible with the collective justification of Rowe's argument. Whether Rowe's argument is sound remains open for debate, as far as my argument here goes, though a debate not centered around skeptical theism anymore.

1. Rowe's Original Evidential Argument

The undeniable existence of evil in our world is the springboard for some of the most powerful arguments for atheism.[1] But there are different ways of formulating such an argument. Offered as a more modest alternative to the "logical" problem of evil—according to which the existence of evil is logically incompatible with the existence of God—William Rowe (1979) is now the *locus classicus* of the "evidential" problem

[1] I will take *atheism* as the belief that an all-good, all-knowing, and all-powerful being, with a personal interest in the lives of his creatures, does not exist. I will refer to such being as a *maximally great God*. I will take *theism* as the belief that a maximally great God exists. Though traditional and still widespread, belief in the existence of some such being is not at all essential to religion and spirituality. The arguments that follow are therefore silent on the rationality of non-traditional (and non-theistic) religious beliefs and practices. See Kaufman (2005) for discussion of the problem of evil in a non-theistic context.

of evil.[2] His argument has two key steps. First, Rowe notes that there are many instances of suffering for which none of us can offer a justifying-reason on God's behalf. We see no outweighing good that could only be furthered or respected by God's allowing for that or worse instances of evil; we see no comparable or worse evil that could only be prevented by God's allowing for that instance of evil. Second, Rowe suggests as implausible that there is something that all of us are missing about each and everyone of these many instances of suffering. How could all of us, for so long, be so blind? In this way, Rowe shifts the focus of the problem of evil from the contradictory existence of any evil whatsoever to the merely unexpected existence of a very specific kind of evil: *pointless suffering*.

We can state Rowe's reasoning more carefully in the following way:

> **Rowe's Original Evidential Argument:**
>
> **1.** After careful reflection on our experience and knowledge, we see no justifying-reason for many instances of suffering in the world.
>
> **2.** If, after careful reflection on our experience and knowledge, we see no justifying-reason for many instances of suffering in the world, then it is reasonable to believe that there is no justifying-reason for at least some of all that suffering.
>
> **3.** So it is reasonable to believe that there is no justifying-reason for at least some of the suffering in the world.
>
> **4.** If there is even one instance of suffering for which there is no justifying-reason, then God does not exist.
>
> **C.** So it is reasonable to believe that God does not exist.[3]

Let me be clear about a central notion. A *justifying-reason* is some consideration in favor of allowing the relevant instance of suffering that would mitigate its existence—furthering or respecting a greater good; preventing a comparable or greater evil. The notion is quite ordinary. Suppose a parent is deciding on whether to get her child

[2]See David Hume (1779, 63) for a classic statement of the logical problem of evil. See Mackie (1955) and Plantinga (1974) for recent and influential discussions.

[3]This is an extended version of the argument Rowe (1979, 336) states explicitly. It includes Rowe's (1979, 337-338) defense of his argument as premises 1 and 2 (which, as Rowe (2006, 81-2) notes, is the proper target of the skeptical theist reply), and it incorporates his later qualifications to the strength of the argument. In further work, Rowe continued to modify and refine his original argument in the face of worries and replies (see, e.g., Rowe 1991, 72-3; 1996, 263-4). My focus in this paper is not, therefore, Rowe's argument in the more abstract sense of "the argument behind Rowe's many explicit articulations." My focus is on Rowe's original reasoning, as stated above.

vaccinated. This will mean allowing for an instance of suffering, but a justifying-reason is readily available: this instance of suffering will significantly lower the chances of much worse suffering in the future. Rowe's argument, in short, exploits the common impression that we are simply unable to produce some such justifying-reason excusing God from much of the suffering in our world.

The heart of Rowe's original argument is premise (2). Yet two of its essential features are often ignored. First, though premise (2) claims that *seeing no* justifying-reason is a good ground for the belief that *there is no* justifying-reason, Rowe is not thereby endorsing the general legitimacy of a "seems so, is so" type of inference. There is a significant amount of epistemic work, that is, being done by Rowe's (1979, 338) qualification: "in the light of our experience and knowledge." We can therefore say that premise (2) captures not a principle of credulity but rather a principle of *responsible* credulity: when in an appropriate epistemic position—acquired by careful reflection on our experience and knowledge—a "seems so, is so" type of inference becomes appropriate as well. Second, premise (2) does not say that seeing no justifying-reason for one instance of suffering makes reasonable the belief that there is no justifying-reason for that same instance of suffering. It says instead that seeing no justifying-reason for *many* instances of suffering is a good ground for the belief that there is no justifying-reason for *at least some* of all that suffering. There is a significant amount of epistemic work, once again, being done by Rowe's (1979, 336) focus on *the amount* of apparently pointless suffering in the world, his argument being "based on *the profusion* of one sort of evil in our world" (namely, the apparently pointless sort). We can therefore say, more accurately, that premise (2) captures not just a principle of responsible credulity but rather a principle of responsible *large-sample* credulity: when in an appropriate epistemic position—acquired by careful reflection on our experience and knowledge—a "seems so, is so" type of inference, going from a large sample to some of its elements, becomes appropriate as well.[4]

Properly understood, premise (2) is thus quite modest. It is consistent, for example, with acknowledging that there might be, for all we know, a justifying-

[4]Notice the progression of Rowe's argument (p. 337) from the 8th paragraph (where he introduces his famous fawn example as an instance of apparently pointless suffering), into the 9th paragraph (where he grants that there might be a justifying-reason behind the fawn example that we just can't see), through the 10th paragraph (where he suggests that we shouldn't look for a proof for the belief in pointless evil but rather for its rational grounds), to the 11th paragraph (where he finally rests his argument on an inference from the large-sample of apparently pointless evil that we see).

reason behind any one particular instance of apparently pointless suffering. It is even consistent with it being unreasonable to believe that there is no justifying-reason for any one particular instance of apparently pointless suffering. As Rowe (1979, 337-8) puts it, taking the intense suffering experienced by a fawn during a forest fire as his example:

> But even if it should somehow be reasonable to believe [that there is a justifying-reason for] the fawn's suffering, we must then ask whether it is reasonable to believe [that there is a justifying-reason for] *all* the instances of seemingly pointless human and animal suffering that occur daily in our world... In the light of our experience and knowledge of the variety and scale of human and animal suffering in our world, the idea that none of this suffering [lacks a justifying-reason] seems an extraordinarily absurd idea, quite beyond our belief.[5]

This is a very clear example of Rowe's principle of responsible large-sample credulity at work, together with a careful acknowledgement of its modesty. While Rowe (2006, 79-80) himself remained resolute that even a single instance of apparently pointless suffering makes it somewhat reasonable to believe that there is no justifying-reason for at least some of the apparently pointless suffering that we see, it is clear that he does not hang his original argument's hat on the evidential support offered by any single instance, taken by itself. Instead, Rowe's argument turns on the much stronger evidential support that is offered by adding together several instances of apparently pointless suffering.

2. The Inductive Justification of Premise (2)

Given the nature of Rowe's premise (2), capturing as it does a principle of responsible large-sample credulity, it is perhaps surprising that much of the discussion that his paper has received—including discussion by Rowe himself—has focused on the appropriateness of the following narrower inference:

> **1*.** We see no justifying-reason for a certain instance of suffering.
>
> **C.** It is likely that there is no justifying-reason for that instance of suffering.

Strictly speaking, such inferences ignore the epistemic work done by both the responsibility qualification and the amount of apparently pointless suffering under

[5] I am translating talk of *outweighing goods* to talk of *justifying-reasons*. The emphasis is original.

consideration. Here are three short examples from three different decades of discussion:

> The form of Rowe's move, I am claiming... is instead from "it appears that not-p" to "it is reasonable to believe that not-p." (Wykstra 1984, 82-3)

> Clearly, [Rowe's argument] depends on an inference from 'so far as I can tell, p' to 'p' or 'probably, p.'" (Alston 1996, 102)

> According to [Rowe], our not seeing any compensating good in cases like these is a prima facie reason to think that there is no compensating good. (McBrayer 2009, 79-80)

From such a sample, one is tempted to conclude that Rowe's original argument has been completely misunderstood.[6] But this is not the case. On one natural way of justifying premise (2), this narrower focus is entirely proper.

On what I will call *the inductive justification* of premise (2), Rowe's argument depends on the cumulative evidential support that is gathered by enumerative induction. Taking for granted that the operative notion of evidential support is best captured by the probability calculus (and that the normative status of "reasonable belief" supervenes on the strength of one's evidence), the picture is something like the following. First, after seeing no justifying-reason for a single instance of suffering (e_1), we have a modicum of evidential support for the belief that there is no justifying-reason for at least some instances of suffering (h). That is:

$$P(h) < P(h \mid e_1)$$

For short, I will say that in such cases e_1 is *evidence* for us that h. Second, after seeing no justifying-reason for a handful of such instances (e_1, e_2, e_3), we have a little bit more support for that same belief. That is:

$$P(h \mid e_1) < P(h \mid e_1 \wedge e_2 \wedge e_3)$$

Third, after seeing no justifying-reason for very many instances of apparently pointless suffering, the belief that there is no justifying-reason for at least some instances of suffering receives enough evidential support to rise up to the status of being reasonable. That is, say:

[6] For more examples of this narrower focus, see Wykstra (1996, 128) and Rowe (2006, 79-80).

$$P(h \mid e_1 \wedge e_2 \wedge e_3 \wedge ...e_n) = .95$$

For short, I will say that in such cases ($e_1 \wedge e_2 \wedge e_3 \wedge ...e_n$), together, is *strong evidence* for us that h.[7] On this inductive justification, then, even though premise (2) licenses as reasonable only a believe that is based on the cumulative evidential support of many instances of apparently pointless suffering—only on such strong evidence—it nonetheless depends on the fact that each of these particular instances carries some evidential weight on its own.

It is only by assuming the inductive justification of premise (2) that we can make sense of the nature of the ensuing discussion of Rowe's argument. Fueled as they are by the inductive justification of Rowe's premise (2), that is, discussions of the evidential problem of evil are commonly centered around the following three claims:

> **(a)** Our seeing no justifying-reason for one instance of suffering *increases* the likelihood for us that there is no justifying-reason for at least some instances of suffering.
>
> **(b)** *Because* of (a), our seeing no justifying-reason for many instances of suffering makes it *very likely* for us that there is no justifying-reason for at least some instances of suffering.
>
> **(c)** If it is very likely for us that p, then we are reasonable in believing that p.

Though premise (2) respects the fact that apparently pointless instances of suffering are by themselves too weak to make reasonable someone's belief that there is no God, it nonetheless claims that, in accumulation, the apparent pointless suffering in the world is strong enough for disbelief. If all three of these claims are true, of course, then premise (2) is true as well. If one of these claims is false, on the other hand, then support for premise (2) is undercut.

3. Skeptical Theism Against Rowe's Induction

Just as its key premise, the conclusion of Rowe's argument is modest as well. It claims that atheism is reasonable, but not that only atheism is reasonable. Rowe (1979, 340-1) himself is quite explicit about his argument being consistent with

[7] Wykstra (1996, 132) calls this *levering evidence*: evidence that is "sufficient to move one from one rational square belief-state to another (given ample but not overwhelming warrant for the initial belief-state)."

theism being reasonable as well. He characterized this position as *friendly atheism*.[8]

Such friendliness, however, has not been returned. The most common reply to Rowe's argument has been that it does not, in fact, show that it is reasonable to be an atheist. The main target for this response has been premise (2).[9] The position licensing a denial of this seemingly modest claim about what's reasonable to believe given a large sample, after careful reflection on our experience and knowledge, combines a claim about the nature of undercutting defeat with a claim about our limited epistemic position. Together, these claims issue the result that our not seeing a justifying-reason for an instance of apparently pointless suffering is not evidence for us that it is not there. Call this resulting position *skeptical theism*.

Thus formulated, skeptical theism is a claim about the *total* defeat of the evidential import of evil: once we appreciate the limits of our epistemic position, we see that instances of apparently pointless evil do not increase the likelihood of there being pointless evil (cf. Bergmann 2008, 387-8; Otte 2012, 127). Sometimes, however, skeptical theism is formulated as a more modest claim about the *partial* defeat of the evidential import of evil: apparently pointless evil may increase the likelihood of there being pointless evil, alright, but never enough to make it reasonable to believe it (cf. Wykstra 1996, 131; Wykstra and Perrine 2012, 384). Yet, given the inductive justification of premise (2), which is required to make sense of the widespread focus on the evidential import of particular instances of apparently pointless evil, the latter formulation of skeptical theism is a mistake. Apparently pointless evil, in this case, is strong evidence for pointless evil by virtue of the *accumulation* of particular instances of non-strong evidence. In order to deny that apparently pointless evil can provide strong evidence for pointless evil, one must deny that it provides any evidence whatsoever in the first place. One cannot grant the evidential power of the particular instances, that is, while denying the increased evidential power of their accumulation.[10]

[8]This turns on the fact that the conditional probabilities required for Rowe's argument depend on plausible, but not rationally required, assumptions about our background knowledge.

[9]Though both sides of this debate have by and large accepted all of the remaining premises, see Hasker (2004), van Inwagen (2006), and McGregor (2012) for surprising denials of premise (4).

[10]In formulating skeptical theism as I do, moreover, I stay faithful to its original ambitions, even if not to some of its later transformations: "Let us say that evidence e weakly supports (or confirms) claim c when e makes c to some degree more likely to be true than it would be on the antecedent evidence... In what follows, I shall argue that the evidence of suffering, as Rowe adduces it, does not disconfirm theism (or confirm atheism) even in the weak sense" (Wykstra 1984, 78-9). This requires denying that "the probability of there being [apparently pointless] evil given atheism is higher than the probability of there being [apparently pointless] evil given theism" (cf. Benton, Hawthorne,

Nonetheless, there are at least two versions of the claims offered in support of skeptical theism, formulated as a claim about the total defeat of the evidential import of evil. One version of these claims is as follows:

> **(expectation$_1$)** If we can expect not to see something, were it to be there, then *our not seeing it* (e) is not evidence for us that *it is not there* (h).

$$[P(e \mid \neg h) = high] \rightarrow \neg[P(h) < P(h \mid e)]$$

> **(limitation$_1$)** We can expect not to see a justifying-reason for instances of apparently pointless suffering, were such reasons to be there.

$$P(e \mid \neg h) = high$$

Here is another version of them:

> **(expectation$_2$)** If we cannot expect to see something, were it to be there, then *our not seeing it* (e) is not evidence for us that *it is not there* (h).

$$[P(\neg e \mid \neg h) = .05] \rightarrow \neg[P(h) < P(h \mid e)]$$

> **(limitation$_2$)** We cannot expect to see a justifying-reason for instances of apparently pointless suffering, were such reasons to be there.

$$P(\neg e \mid \neg h) = .05$$

Both versions of (expectation) claim that certain facts about our epistemic position defeat the evidential support that would otherwise be offered by a certain experience. (expectation$_1$), however, states a weaker condition than (expectation$_2$): (expectation$_2$) claims that evidential support is undercut not only when one *actually has* a certain expectation not to see x, but also when one *cannot have* a certain expectation to see x. Any case that fails (expectation$_2$), therefore, also fails (expectation$_1$). Both versions of (limitation), in turn, make claims about the limitations of our epistemic position. (limitation$_1$), however, makes a stronger claim than (limitation$_2$): (limitation$_1$) says that we not only cannot expect to see God's justifying-reasons, we can positively expect not to see them at all. If (limitation$_1$) is true of us, that is, then (limitation$_2$) is true as well. Unfortunately, proponents and

& Isaacs 2016, 5), and skeptical theists have not been shy about doing so (cf. Howard-Snyder & Bergmann (2004, 22).

critics of skeptical theism have not always been clear about which combination of these claims they find plausible or implausible.[11] For present purposes, I will assume that skeptical theism is supported by the truth of the stronger (expectation$_2$) and the weaker (limitation$_2$).

We can illustrate the appeal of these claims with some examples. Suppose you are wondering whether a certain hypodermic needle, about to go into your arm, is contaminated (cf. Wykstra 1996, 129; 2007, 88). Suppose your doctor, in an effort to comfort you, fixes her stare at the needle while saying "I see no viruses on this needle." Is her experience any evidence that the needle is virus-free? Does her seeing no viruses on the needle increase the likelihood for her that there are no viruses on the needle? Not only is the answer to these questions 'no', (expectation$_2$) seems to be their explanation. It is precisely because the doctor cannot expect to see any viruses, were there any on the needle, that seeing no virus upon inspection falls short of being evidence against that claim. Now consider once again the parents who allow their child one bit of suffering on the basis of the justifying-reason of preventing much worse suffering in the future (cf. Wykstra 1984, 88; 1996, 130). Given the epistemic position of the child, in relation to the epistemic position of her parents, she cannot expect to see the justifying-reason on the basis of which her parents are allowing her that mystifying bit of suffering. (limitation$_2$), then, is simply the plausible claim that our epistemic position in relation to God's is similar to the child's epistemic position in relation to her parents'. Just as the child, we cannot expect to see God's justifying-reasons for allowing the many apparently pointless bits of suffering that we see, were any such reasons to truly be there.

If these supporting claims are true, however, then (a) is false: our seeing no justifying-reason for one instance of suffering *does not* increase the likelihood for us that there is no justifying-reason for at least some instances of suffering. An experience of apparently pointless suffering, that is, is no evidence for there being actual instances of pointless suffering. According to (expectation$_2$), after all, our seeing no justifying-reason can only be evidence if we can expect to see it, were it to be there. In Rowe's (2006, 88) words, if we are to accept (a), then we "must

[11] For endorsements of the weaker (expectation$_1$), see Rowe (1984, 96) and Kraay (2007, 210); for endorsements of the stronger (expectation$_2$), see Wykstra (1984, 85; 1988, 143; 1996, 133), Plantinga (2000, 466), Perrine and Wykstra (2012, 377). For endorsements of the stronger (limitation$_1$), see Wykstra (1984, 87-89; 1988, 143), Plantinga (2000, 467), Kraay (2007, 211-15), McBrayer (2010), Wykstra and Perrine (2012, 375), DePoe (2014, 42-4); for endorsements of the weaker (limitation$_2$), see Wykstra (1996, 135-7), Bergmann (2001, 279), Howard-Snyder (2009, 18), Long (2014, 65-66).

suppose that the goods we know of are representative of the goods there are." But (limitation$_2$) tells us that we cannot expect to see some such reason, were it to be there; it tells us that we cannot suppose that the goods we know of are representative of the goods there are. So our seeing no justifying-reason for a particular instance of apparently pointless suffering does not increase the likelihood of there not being any such reason for at least some instances of suffering. Skeptical theism is true, then, and (a) is false.[12]

According to its inductive justification, recall, Rowe's premise (2) depends on some evidential support being offered by individual instances of apparently pointless suffering. Yet, if skeptical theism is correct, each individual instance offers no support at all. Seeing no justifying-reason for more and more cases of suffering, then, does not push us closer and closer to the threshold of evidential support required for reasonable belief. In this way, skeptical theism seems to undercut the support adduced for Rowe's premise (2). Though Rowe's argument aims to show merely that it is reasonable to be an atheist—on the basis of the strong evidence provided by adding together very many instances of apparently pointless suffering—skeptical theism seems to show that Rowe's argument fails even in the modest task that was its charge.

4. The Evidential Dimension of the Preface Paradox

Leaving aside debates over skeptical theism, in the rest of this paper I will argue that there is a justification available for premise (2) that does not depend on (a):

> **(a)** Our seeing no justifying-reason for one instance of suffering *increases* the likelihood for us that there is no justifying-reason for at least some instances of suffering.

On what I will call *the collective justification* of premise (2), Rowe's argument does not depend on there being some modicum of evidential support provided by particular instances of apparently pointless suffering. The key to this justification is seeing

[12] I do not mean to suggest that this is obviously so. For resistance to the stronger (limitation$_1$), see Rowe (1984, 96, 2006, 85-88); for resistance to the weaker (limitation$_2$) see Rowe (1991, 73; 1996, 272; 2006, 88-91). For resistance to (limitation) in general, see Dougherty (2012, 21-23). For resistance to (expectation$_2$), see Benton, Hawthorne, & Isaacs (2016, 15-16). For the claim that (expectation) in general is in tension with commonsense epistemology, see Dougherty (2008, 176). Such formidable *résistance* notwithstanding, I am here treating skeptical theism as successful since the dialectical force of my discussion consists in the very suggestion that Rowe's original argument does not depend on the success of these replies.

that the difference between the epistemic status of our judgments about pointless evil given a particular instance of apparently pointless evil and the epistemic status of that judgment given a collection of these instances needs not supervene on a continuum of increased evidential support. Instead, the relation between the epistemic status of our judgments in these two cases can instead reflect a paradoxical evidential feature of certain inductive bases that is often ignored. I identify and clarify this general feature in this section, and I apply it to the problem of evil in the next.

Consider what is known as the preface paradox.[13] The preface paradox arises in cases where we reasonably believe *of each claim* in a collection that it is true and yet are unreasonable in believing *of the collection* that it does not contain a claim that is false. A book provides a common instance of these kinds of cases: though the careful and responsible writer of the book reasonably believes every claim that she makes, she would exceed in hubris if she also believed that she has made no mistake throughout. Hence the common confession found in prefaces, after acknowledgement of the positive influence of friends and peers, that "any mistakes are wholly mine." Common as they are, these cases are paradoxical since they seem to violate a very plausible principle of *aggregation*:

> **The Preface Paradox**
>
> **PP$_1$.** I reasonably believe that p is true, I reasonably believe that q is true, ... I reasonably believe that x is true.
>
> **PP$_2$.** I am not reasonable in believing that (p & q & ... x) is true.
>
> **PP$_3$.** If I reasonably believe that p is true, and I reasonably believe that q is true, then it is reasonable for me to believe that (p & q) is true as well.

These three claims cannot all be true. Yet we seem to have instances—choose any carefully written book—were all three seem true. That's the preface paradox.

Few attempts to solve this paradox take aim at either (PP$_1$) or (PP$_2$).[14] It seems exceedingly plausible that these are true, and turning to talk of degrees of belief is a common way of making a case for them. Suppose I am extremely confident in each of the claims I make in my book; I'm 95% confident that p, and 95% confident that q, and so on. Given the plausible assumptions already mentioned in section 2, this means that I am reasonable in believing that each of them is true. Hence (PP$_1$).

[13] See Makinson (1965) for an early treatment; see Christensen (2004) for a more recent discussion.

[14] See, however, Pollock (1990), Ryan (1991), Kaplan (1996), and Adler (2002) for arguments against the rational conjunction of (PP$_1$) and (PP$_2$). See also Christensen (2004, 33-68) for a critical reply.

Yet no matter how confident I am in each of my claims individually, the probability calculus tells me that their conjunction remains wildly unlikely. At 95% confidence, we need to conjoin only 15 claims before we get that their likelihood, together, is very low: $(p \ \& \ q \ \& \ ... \ \ x) = .46$. From this it follows that I am not reasonable in believing that $(p \ \& \ q \ \& \ ... \ \ x)$ is true. Hence (PP_2). So it seems we have good reason to resolve the preface paradox by rejecting (PP_3).

There is, however, an evidential dimension to the preface paradox that is often ignored and which is resilient even under the probabilistic resolution sketched above. We can see this by noticing that it is false that an apparently true claim (p) is *evidence* for me that there is at least one false claim in the book $(\neg(p \ \& \ q \ \& \ ... \ \ x)$, or h for short). That is:

$$\neg[P(h) < P(h \mid p)]$$

It is just not the case that, upon carefully considering the content of one of my claims and upon renewing my confidence that it is true, my book then becomes more likely for me to have at least a false claim in it. In fact, the exact opposite seems true. Upon carefully considering the content of one of my claims and upon renewing my confidence that it is true, it seems my book should become less likely for me to have at least a false claim in it. Think about it: suppose I want to check whether there is at least one false claim in the book (h); the first claim I check is apparently true; shouldn't I then lose a little confidence in h? I certainly shouldn't gain any confidence in it; I shouldn't say "this claim seems true; So I'm now a bit more confident that there is something false in this book". This is an intuitive assessment of the evidential impact of an apparently true claim on my judgments about the book that must be respected by our understanding of the preface paradox.[15] Nonetheless, it seems perfectly correct to say that when I consider the collection of all my claims I therefore have evidence—strong evidence, even—for h. So while an apparently true claim in my book, considered individually, is no evidence for me that there is at least one false sentence in my book, a large collection of these claims is, in fact, strong evidence for that same belief. This is the evidential dimension of the preface paradox. What it reveals is that the fact that it is not reasonable for me to believe that $(p$

[15]Notice that while modeling the *evidence for* relation on the probability calculus can give us a deeper understanding of the behavior of that relation, it cannot directly tell us what is evidence for what all the way down. First we tell the model enough fundamental facts about what we take as evidence for what; then the model tells us the full implications of that assignment.

& q & ... x) is true is not a consequence of the accumulation of *evidence* provided by p and q, and so on, by enumerative induction. In other words, the evidential dimension of the preface paradox reveals it as a case where the difference between the epistemic status of our judgment about truth given a particular proposition and the epistemic status of that judgment given a collection of these propositions is not supervening on a continuum of increased evidential support. This dimension of the paradox, notice, survives the probabilistic resolution of the preface paradox via rejection of (PP_3).

What characterizes this paradoxical relationship between the epistemic status of our judgment given particular instances and the epistemic status of our judgment given a large collection of those instances is the presence of what I will call a *reversing condition*. In the case of the preface paradox, the reversing condition seems to be the *modest fallibility* of a careful and responsible writer. Since we are not hopelessly fallible, we are reasonable in believing that claims that we are confident are true, after careful consideration, are really true. But since we are fallible nonetheless, we are not reasonable in believing that a large amount of these claims—as found in a book—are all true. This is not to say that this reversing condition explains away the mystery behind the evidential dimension of the preface paradox. The mismatch between our intuitive assessment of the evidential impact of apparently true claims and the impact of their collection remains in place. But identifying a relevant reversing condition is perhaps a way to put a finger on an important part of the underlying mechanism.

5. The Collective Justification of Rowe's Premise (2)

I now want to suggest that something analogous could be true of the evidential problem of evil. Even though we are reasonable in believing *of each instance* of apparently pointless suffering that it is not in fact a case of pointless suffering, we are nonetheless not reasonable in believing *of the collection* of these instances that it does not contain at least one instance of pointless suffering. Call this the *the paradox of evil*:

> **The Paradox of Evil**:
>
> **PE_1.** I reasonably believe that instance of suffering a has a point, I reasonably believe that instance of suffering b has a point, ... I reasonably believe that instance of suffering x has a point.

PE$_2$. I am not reasonable in believing that (a & b & ... x) has a point.

PE$_3$. If I reasonably believe that a has a point, and I reasonably believe that b has a point, then it is reasonable for me to believe that (a & b) has a point as well.

What we have here is a case where the difference between the epistemic status of our judgment about justifying-reasons given a particular instance of apparently pointless suffering and the epistemic status of that judgment given a collection of these instances is not supervening on a continuum of increased evidential support. We have seen in our discussion of skeptical theism that it is false that an instance of apparently pointless suffering (e) is evidence for me that there is at least one instance of apparently pointless suffering in the world (h). (Recall that I am taking that for granted.) It is just not the case that, upon carefully considering an instance of apparently pointless suffering and upon renewing my confidence that it nonetheless has a point, the world then becomes more likely for me to have at least one instance of pointless suffering in it. That is:

$$\neg[P(h) < P(h \mid e)]$$

But we have also seen in our discussion of the preface paradox that, despite such an assessment of the evidential impact of e on h, it can remain the case that I am not reasonable in believing that h is false given a large collection of instances of apparently pointless suffering. This would mean that the fact that it is not reasonable for me to believe that (a & b & ... x) has a point is not a consequence of the accumulation of evidence provided by a and b, and so on, by enumerative induction. As in the preface paradox, the behavior of the relevant evidence in the problem of evil is a bit of a mystery.

This time, however, what characterizes this paradoxical relationship between the epistemic status of our judgment given particular instances and the epistemic status of our judgment given a large collection of those instances is the presence of a different reversing condition: *God's maximal greatness*—the fact that He is an all-good, all-knowing, and all-powerful being, with a personal interest in the lives of His creatures. Though His maximal greatness prevents the expectation that we would see some justifying-reason for any particular instance of apparently pointless suffering (His capacity for instrumental reasoning being so far beyond ours'), His greatness produces the following disjunctive expectation nonetheless: either His

justifying-reasons for so much apparently pointless suffering would be available to us by now in broad outline, or some intelligible indication of the existence of such justifying-reasons would be available to us as proxy, or His comforting presence would be unmistakably known to us in compensation (more on these in the next section). Just as before, this is not to say that this reversing condition explains away the mystery behind the paradox of evil. But identifying a relevant reversing condition is perhaps a way to put a finger on an important part of the underlying mechanism.

In this way, the collective justification of premise (2) does not depend on (a)—on the claim that our seeing no justifying-reason for one instance of suffering increases the likelihood for us that there is no such reason for at least some instances of suffering. If fueled instead by the collective justification of Rowe's premise (2), that is, discussions of the evidential problem of evil would then be centered around the following three claims:

> (¬ a) Our seeing no justifying-reason for one instance of suffering *does not* increase the likelihood for us that there is no justifying-reason for that instance of suffering.
>
> (b*) *Despite* (¬ a), our seeing no justifying-reason for *many* instances of suffering makes it *very likely* for us that there is no justifying-reason for at least some instance of suffering.
>
> (c) If it is very likely for us that p, then we are reasonable in believing that p.

This, I suggest, is one way to capture "the intuition that the existence of so much suffering for which we are unable to see any point at all disconfirms theism" (Rowe 1984, 100). And if all three of these claims are true, then premise (2) is true as well. In fact, on its collective justification, premise (2) turns out to be even more modest than before: it not only respects the fact that apparently pointless instances of suffering are by themselves too weak to make reasonable someone's belief that there is no God, it also respects the skeptical theist's insistence that these instances, individually, offer no support at all for that belief. It just turns out that instances of evil bear witness against God, we can borrow from Quine in saying, not individually but rather as a corporate body.[16]

[16]This reading is not available for Rowe's (1996, 263-4) later formulation of his argument. There, the amount of evil is irrelevant. What is relevant instead is the amount of justifying-reasons that we can survey and the fact that none of them seems to justify some particular instance of evil under consideration. The relevant inductive inference, in this case, goes from "No justifying-reason

6. Skeptical Theism and The Paradox of Evil

We have seen how skeptical theism can make a plausible case against the claim that particular instances of apparently pointless suffering are evidence for us that there truly is pointless suffering in the world. Borrowing a term from Steve Wykstra—speaking for the character *Bea* in Russell & Wykstra (1988, 143)—the justifying-reason for an instance of suffering has very low *seeability* for us: if the perfect God exists, then we cannot expect to see His justifying-reasons in the first place. But we have also seen that there is a way of justifying Rowe's premise (2) where it does not depend on the very high seeability of these particular justifying-reasons. Borrowing a term from Bruce Russell—speaking for the character *Athea* in Russell & Wykstra (1988, 147)—we can say that the problem of evil does not depend on whether our world is *morally transparent*.

Recall that (limitation$_2$)—the claim that we cannot expect to see a justifying-reason for the instances of apparently pointless suffering, were such reasons to be there—turns on the plausibility of a gap between our epistemic position and God's. Just as a child cannot expect to be able to grasp the justifying-reason behind her parents' actions, we as well cannot expect to be able to grasp the justifying-reasons behind God's. In both cases, what blocks the expectation is a difference in how much is known about the relevant facts by both parties, which in turn is the result of either a temporary lack of access to certain kinds of information or an inherent incapacity to grasp them. As Wykstra (1996, 140) puts it

> If there is a being who created and sustained this universe around us, the wisdom and vision of this being would be considerably greater than our own. Given what we independently know of our *cognitive limits*... the vision of such a being might well be to ours, as a parent's is to that of a one-month-old human infant. (My emphasis)

It is because of our cognitive limits, that is, that we cannot expect to see God's justifying-reasons for any particular instance of apparently pointless suffering. Some of His reasons just may be beyond our ken. But the paradox of evil does not depend on our privileged epistemic access or on our inherent capacities as knowers of the relevant facts. The paradox of evil, that is, does not depend on our warrant for

we know of justifies God in permitting this evil" to "No justifying-reason at all justifies God in permitting evil." This inference is pitted against skeptical theism in ways similar to those discussed in section 3.

believing that "the possible goods we know of are representative of the possible goods there are," (Bergmann 2001, 279).[17]

The paradox of evil turns instead on what we can expect of a maximally great God given the helpless position we find ourselves in on account of our cognitive limitations. We cannot expect to divine his ways; but what can we expect him to do for us about it? The paradox of evil gets off the ground with only modest expectations. Perhaps it is inappropriate to expect a maximally great God to explain to us the precise point of every single instance of apparently pointless suffering. Perhaps it is also inappropriate to expect Him to provide explanations for any large amount of suffering on a timely fashion, or to any one individual in particular. Yet it almost seems to behove Him to explain to us, at some point and in broad outline, His justifying-reasons for allowing for *so many* instances of apparently pointless suffering for so long. At a minimum, it seems to behove Him to either give us some intelligible indication of the existence of such justifying-reasons as proxy—some form of assurance that there are, in fact, justifying-reasons—or to make His comforting presence unmistakably known to us in compensation. Explanations in broad outline; intelligible indications of them as proxy; comforting presence in compensation—none of this depends in any way on our privileged epistemic access or on our inherent capacities as knowers of the relevant facts.[18]

The paradox of evil, in fact, is not only compatible with the skeptical theist suggestion that our epistemic position is to God's like that of a small child to her parents', it is perhaps buttressed by that very analogy. When a loving and capable parent allows her child to undergo a painful procedure—vaccination, surgery, and so on—that parent does all she can to stand beside the child and to assure her that there is, in fact, a point to it all. In those cases, as Rowe (1996, 277) puts it "the parent attends directly to the child throughout its period of suffering, comforts the child to the best of her ability, expresses her concern and love for the child in ways that are unmistakably clear to the child, assures the child that the suffering will end, and tries to explain, as best she can, why it is necessary for her to permit the suffering even though it is in her power to prevent it." Of course, perhaps it is inappropriate to expect a loving and capable father to explain the precise point of

[17]See Dougherty (2012, 21-23) for a challenge to Wykstra's claim that we cannot expect to see God's justifying-reasons "given what we independently know of our cognitive limits."

[18]Notice that the paradox of evil does not require any of this "appearing obvious" (Butler 1833, 266) or having an "irresistible proof" (Paley 1879, 470). See DePoe (2014, 37) for the claim that such extra requirements would undermine the attainment of other goods.

every single instance of apparently pointless suffering that a child undergoes during a painful surgery. Perhaps it is also inappropriate for the child to expect explanations on a timely fashion. Yet it almost seems to behove a loving and capable father to explain to her child, at some point and in broad outline, her justifying-reasons for allowing for so many instances of apparently pointless suffering. At a minimum, once again, it seems to behove her to either give her child some intelligible indication of the existence of such justifying-reasons as proxy—some form of assurance that there are, in fact, justifying-reasons—or to make her comforting presence unmistakably known to her child in compensation. In both cases—the child in relation to her parents and us in relation to God—temporal lack of access to information and inherent limitations on one's cognitive capacities are irrelevant for the reasonableness of the expectation.

The reasonableness of this expectation, notice, depends crucially on its disjunctive nature: either God's justifying-reasons for so much suffering for so long must be available to us at least in broad outline, or some intelligible indication of the existence of such justifying-reasons must be available to us as proxy, or His comforting presence would be unmistakably known to us in compensation. This three-pronged disjunctive expectation deflects worries such as Long's (2014, 68) analogical counterexample:

> There was a time when I allowed my young son to try to walk, knowing that this almost certainly would result in numerous falls, which would sometimes produce pain and frustration. Nevertheless, my allowing my son to try to walk was both justified for me and expressive of my love for him. But, when he was learning, he was in no cognitive position to understand anything close to my reasons for allowing him to try to walk.

Long intends this an example of how suffering can be allowed by a loving parent without a corresponding explanation of his or her justifying-reasons. Crucially, the example exculpates the parent on account of the child's complete lack of cognitive capacity for understanding even the very concept of a justifying-reason. There are far reaching implications for the suggestion that we are cognitively compared to God as a toddler is to its parents (as opposed to as a young child is to its parents); how could we then claim to understand *anything* about God and his plans and desires, even when He attempts to reveal Himself to us? That aside, however, allowing your toddler the pain and frustration of the inevitable falls without making your

comforting presence unmistakably known in compensation is a sure sign of a parent who is falling short of what can be reasonably expected out of love. The same seems true of a maximally great God.

This could give us good reason to think that despite the plausible gap between our epistemic position and God's—captured by (limitation$_2$)—and despite the manner in which this gap plausibly defeats the evidential import of particular instances of apparently pointless suffering—captured by (expectation$_2$)—our seeing no justifying-reason for *many* instances of suffering makes it *very likely* for us that there is no justifying-reason for at least some instances of suffering. Skeptical theism, at any rate, grounded as it is on the cognitive-gap evident in the parent-child analogy, seems incapable of challenging the required three-pronged disjunctive expectation.

Conclusion

Most of the discussion generated by Rowe's justly famous argument has focused on the appropriateness of the inference from "I see no justifying-reason for this instance of suffering" to "there is no justifying-reason for at least some instances of suffering." This focus, however, is only appropriate given what I have called the inductive justification of Rowe's premise (2). Granting the skeptical theist criticism of this inference, I've argued that there is an alternative justification of premise (2) that does not depend on it. According to the collective justification, Rowe's argument, and his friendly atheism, are best understood as grounded on the paradoxical evidential import of the collection of instances of apparently pointless suffering, taken as a corporate body, and not at all on the evidential import of particular instances themselves.

Thus understood, the soundness of Rowe's original argument is compatible with skeptical theism by virtue of (a) respecting the evidential impact of our cognitive limitations, and (b) depending instead on the following disjunctive expectation: either God's justifying-reasons for so much suffering for so long must be available to us in broad outline, or some intelligible indication of the existence of such justifying-reasons must be available to us as proxy, or his comforting presence must be unmistakably known to us in compensation. Importantly, I have not argued that this three-pronged expectation is frustrated. My aim was simply that of showing how skeptical theism and the evidential problem of evil need not be in tension in the first place.

Notice, however, that the collective justification of premise (2), based on the paradox of evil, reveals an important connection between Rowe's argument and *the problem of divine hiddenness*: the claim that God's absence is evidence against his existence.[19] Rowe (2006, 88) himself, in fact, seems to recognize some connection between the two, even if not exactly the connection that I am suggesting here:

> My own inclination is to think that given the horrendous evils in our world, the absence of the God who supposedly walked with Adam and Eve in the garden is evidence that there is no God. For surely, if there were a God he would wish to provide us with strong reasons to think that he exists, given that the horrendous evils in our world, both natural and moral, seem to provide us with reason to doubt his existence.

Perhaps Rowe's suggestion is that, *on top of* the evidence against the existence of God that we get from instances of apparently pointless suffering, the absence of God becomes conspicuous since it (and perhaps it alone) could outweigh the evidential scales in His favor. I make no claim to an accurate exegesis of Rowe's own original intentions with his early argument and with this later passage. But we can understand both that argument and this passage in ways that are compatible with the collective justification of premise (2). It is precisely *because of* the absence of God that the of amount horrendous evils in our world, collectively but not individually, becomes evidence against His existence.

[19] As McBrayer (2010, 612) puts it, "the argument from divine hiddenness claims that if God existed, his loving nature and desire to be in communion with capable creatures would compel him to set up the world so that everyone would believe in him. Thus, God would seek to make his existence obvious. But since it's not obvious that God exists, this fact thereby provides evidence for his non-existence." Dougherty (2012, 21) also notes the "the important connection between the problem of evil and the problem of divine hiddenness." See Schellenberg (1993; 2005a; 2005b) for an influential discussion of divine hiddenness.

References

Adler, Jonathan (2002) *Belief's Own Ethics.* Cambridge: MIT Press.

Alston, William (1996) "The Inductive Argument from Evil and the Human Cognitive Condition," in Daniel Howard-Snyder (ed.), *The Evidential Argument From Evil.* Indiana University Press: 97-126

Benton, Matthew A., Hawthorne, John, & Isaacs, Yoaav (2016). "Evil and Evidence," in Jonathan Kvanvig (ed.) *Oxford Studies in Philosophy of Religion* 7: 1-31.

Bergmann, Michael (2001) "Skeptical Theism and Rowe's New Evidential Argument From Evil", *Nous* 35(2): 278-296.

Bergmann, Michael (2008). "Skeptical Theism and the Problem of Evil," in Thomas P. Flint & Michael C. Rea (eds.), *The Oxford Handbook of Philosophical Theology.* Oxford University Press: 374-399.

Butler, Joseph (1833). *The Analogy of Religion.* Oxford University Press.

Christensen, David (2004) *Putting Logic in Its Place.* Oxford University Press.

DePoe, John M. (2014) "On the Epistemological Framework for Skeptical Theism," in Trent Dougherty and Justin P. McBrayer (eds.), *Skeptical Theism: New Essays.* Oxford University Press: 32-44.

Dougherty, Trent (2008). "Epistemological Considerations Concerning Skeptical Theism". *Faith and Philosophy* 25(2): 172-176.

Dougherty, Trent (2012). "Reconsidering the Parent Analogy: Unfinished Business for Skeptical Theists". *International Journal for Philosophy of Religion* 72(1): 17-25.

Hasker, William (2004). *Providence, Evil, and the Openness of God.* Routledge.

Howard-Snyder, Daniel (2009) "Epistemic Humility, Arguments From Evil, and Moral Skepticism." *Oxford Studies in Philosophy of Religion* 2: 17-57.

Howard-Snyder, Daniel & 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.* Blackwell Publishers: 13-24.

Hume, David (1779/1980) *Dialogues concerning Natural Religion.* Indianapolis: Hackett Publishing.

Kraay, Klaas J. (2007) "Absence of Evidence and Evidence of Absence" *Faith and Philosophy* 24: 203-228.

Kaplan, Mark (1996) *Decision Theory as Philosophy* New York: Cambridge University Press.

Kaufman, Whitley (2005) "Karma, Rebirth, and the Problem of Evil," *Philosophy East and West* 55(1): 15-32.

Long, Todd R. (2014) "Minimal Skeptical Theism," in Trent Dougherty and Justin P. McBrayer (eds.), *Skeptical Theism: New Essays*. Oxford University Press: 63-75.

MacGregor, Kirk R. (2012) "The Existence and Irrelevance of Gratuitous Evil" *Philosophia Christi* 14(1): 165-80.

Makinson, D.C. (1965) "The Paradox of the Preface," *Analysis* 25: 205-7.

Mackie, John (1955) "Evil and Omnipotence," *Mind* 64(254): 200-212

McBrayer, Justin P. (2009) "Cornea and Inductive Evidence," *Faith and Philosophy* 26(1):77-86.

McBrayer, Justin P. (2010) "Skeptical Theism," *Philosophy Compass* 5(7): 611-623.

Otte, Richard. (2012). "Comparative Confirmation and the Problem of Evil," in Jake Chandler and Victoria S. Harrison (eds.), *Probability in the Philosophy of Religion*. Oxford: Oxford University Press: 127-143.

Paley, William (1879). *A View of the Evidences of Christianity*. Robert Carter & Brothers.

Plantinga, Alvin (1974) *God, Freedom, and Evil*. New York: Harper and Row.

Plantinga, Alvin (2000) *Warranted Christian Belief*. Oxford University Press.

Pollock, John (1990) *Nomic Probability and the Foundations of Induction*. Oxford University Press.

Rowe, William (1979) "The Problem of Evil and Varieties of Atheism," *American Philosophical Quarterly* 16(4): 335-341.

Rowe, William (1984) "Evil and the Theistic Hypothesis: A Response to Wykstra," *International Journal for Philosophy of Religion* 16(2): 95-100.

Rowe, William (1991) "Ruminations about Evil," *Philosophical Perspectives* 5: 69-88.

Rowe, William (1996) "The Evidential Argument From Evil: A Second Look," in Daniel Howard-Snyder (ed.), *The Evidential Argument From Evil*. Indiana University Press: 262-85.

Rowe, William (2006) "Friendly Atheism, Skeptical Theism, and the Problem of Evil," *International Journal for Philosophy of Religion* 59(2): 79-92.

Ryan, Sharon. (1991) "The Preface Paradox," *Philosophical Studies* 64: 293-307.

Schellenberg, J.L. (1993) *Divine Hiddenness and Human Reason*. Cornell University Press.

Schellenberg, J.L. (2005a) "The Hiddenness Argument Revisited (I)," *Religious Studies* 41(2): 201-215.

Schellenberg, J.L. (2005b) "The Hiddenness Argument Revisited (II)," *Religious Studies* 41(3): 287-303.

van Inwagen, Peter (2006) *The Problem of Evil*. Oxford University Press.

Wykstra, Stephen (1984) "The Humean Obstacle to Evidential Arguments from Suffering: On Avoiding the Evils of 'Appearance'" in *International Journal for the Philosophy of Religion* 16: 73-93.

Wykstra, Stephen and Russell, Bruce (1988) "The Inductive Argument from Evil: A Dialogue," *Philosophical Topics*, 16(2): 133-160.

Wykstra, Stephen (1996) "Rowe's Noseeum Argument from Evil," in D. Howard-Snyder (ed.), *The Evidential Argument from Evil*. University of Indiana Press: 127-151

Wykstra, Stephen (2007) "CORNEA, Carnap, and Current Closure Befuddlement," *Faith and Philosophy* 24(1): 87-98.

Wykstra, Stephen and Perrine, Timothy (2012) "Foundations of Skeptical Theism: CORNEA, Counterfactuals, and Conditional Probabilities," *Faith and Philosophy* 29(4): 375-399.