

# Why Context Matters for the Problem of Evil

Ian M. Church, Justin L. Barrett, Blake McAllister, and James S Spiegel.

Recent empirical research on Rowe's influential 1979 formulation of the problem of evil has found that adding broader background information to Rowe's famous fawn vignette significantly diminishes participants' perception that the suffering in the case he describes is pointless (see Church et al. 2021). Call this *the context effect*. This finding raises an important question: what is it about context that has this effect? Does context diminish perceived pointlessness because it provides a possible point or explanation for the suffering? Does it only occur when context is positively valenced? Is it because the context distracts participants from the suffering? Or is it because the context reveals the limits of our epistemic position? These are the central questions this paper aims to address.

In §1, we review the recent empirical literature highlighting the role of context in judgments about Rowe-style suffering. In §2, we identify several hypotheses about why context might diminish perceived pointlessness. In §3, we describe a follow-up study designed to test these hypotheses. In §4, we present our results. In §5, we discuss their philosophical significance and consider whether the effect of context may be connected to the broader role of narrative in human cognition.

Our results suggest that none of the most natural explanations fully accounts for the context effect. The central lesson is not that context defeats Rowe. It is that context complicates the evidential role of Rowe-style intuitions. Context appears to play an important role in diminishing perceived pointlessness, but it does not seem to do so merely because: 1) it provides an explicit theodicy, 2) it is positively valenced, 3) it distracts participants from the target suffering, or 4) it prompts skeptical-theist reflection on human cognitive limitations. The role context plays in judgments about suffering therefore remains partly mysterious. We close by suggesting that this mystery may point toward a broader connection between context, narrative, and the human tendency to understand suffering by embedding it within larger structures of intelligibility.

## §1: Background

In William Rowe's seminal version of the problem of evil,<sup>1</sup> he levels the following argument against theism:

1. There exist instances of intense suffering which an omnipotent, omniscient being could have prevented without thereby losing some greater good or permitting some evil equally bad or worse.
2. An omniscient, wholly good being would prevent the occurrence of any intense suffering it could, unless it could not do so without thereby losing some greater good or permitting some evil equally bad or worse.
3. [Therefore,] there does not exist an omnipotent, omniscient, wholly good being. (1979, 336)

---

<sup>1</sup> Modified versions of Rowe's argument can, of course, be found in later work—see for example, Rowe, 1996—but empirical work has primarily focused on Rowe's 1979 formulation because it is the most influential and the most cited variation of the problem in the academic literature.

Of course, such an argument is valid, but why should we think that the premises are true? Premise 2 seems fairly unobjectionable. Indeed, as Rowe notes, “This premise...is, I think, held in common by many atheists and nontheists” (1979, 336). For this paper, we’re happy to agree; our focus will be on premise 1, which identifies what we, with Rowe, are calling *pointless suffering*.

Why should we think that there is pointless suffering? Here Rowe has us think about what seems like a good candidate:

**Fawn:** Suppose in some distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering. (1979, 337)

According to Rowe, “so far as we can see, the fawn’s intense suffering is pointless” (1979, 337). (Though, whether or not the suffering is *actually* pointless is the subject of debate.) While an omnipotent, omniscient, all-good being certainly could have prevented such an event, it’s extremely difficult to imagine how permitting something like the suffering of FAWN could either prevent a greater evil from occurring or usher in some greater good. As such, premise 1 looks plausible.

But, as Rowe is quick to note, this doesn’t amount to a proof. For all we can tell, there *is* a greater evil or good that allowing FAWN prevents or affords, respectively. The problem, as Rowe sees it, is that given “*our* experience and knowledge of the variety and profusion of suffering in our world” it sure seems like such evils are pointless; hence, the argument still provides “rational support for atheism” and makes it “reasonable for us to believe that the theistic God does not exist” (1979, 338, emphasis ours).<sup>2</sup>

Critically, it is our intuitions regarding FAWN (and related cases) that are the driving force for thinking that premise 1 is true and, thus, for Rowe’s evidential argument more generally. As Alvin Plantinga elucidates Rowe’s argument, if it *seems* as though the suffering in FAWN is pointless, then that gives us a reason for thinking that the suffering in FAWN is pointless (2000, 465–66).

Recent experimental work has investigated whether ordinary judgments about Rowe’s fawn case are sensitive to features of the way the case is presented. In McAllister et al. 2024, the authors report a study designed to examine both the distribution of Rowe-style intuitions across demographic groups and the psychological mechanisms that might underwrite such intuitions. The study tested several possible influences on judgments about the fawn’s suffering, including participants’ demographic characteristics, the animal described in the vignette, the presence or absence of a picture, and the presence or absence of broader background information or context.

For present purposes, the most important of these variables is context. The fawn case is strikingly brief: in Rowe’s original presentation, the relevant vignette is only two sentences long. This sparseness may matter. Some philosophers have stressed the importance of context and narrative in addressing the problem of suffering.<sup>3</sup> One might therefore wonder whether the appearance of pointlessness in Rowe’s fawn case is partly a function of the way the case is presented: isolated from broader causal, ecological, temporal, or narrative background. Would the inclusion of additional background information alter participants’ perception that the suffering is pointless?

To investigate this and related questions, Church, Carlson, and Barrett (2020) and Church, Warchol, and Barrett (2022) reported an empirical study, the results of which were more fully reported in McAllister et al. 2024.<sup>4</sup> Participants read some variation of Rowe’s fawn vignette. For

---

<sup>2</sup> That said, it is somewhat unclear precisely how much evidential weight Rowe ascribes to his argument. For a fuller consideration of the array of possible interpretations, see Wykstra 1996.

<sup>3</sup> See, for example, Stump 2010.

<sup>4</sup> The experiment was pre-registered on the Open Science Framework here: <https://osf.io/ebgpd>

half of the participants, the vignette was accompanied by broader background information about forest fires, including their role in clearing dead organic material, enriching soil, enabling new growth, and reducing the likelihood of more catastrophic fires. The other half of participants read the vignette without this background information, just as it appeared in Rowe's 1979 paper. The study also varied the animal described in the vignette, using a fawn, boar, or vulture, and varied whether participants were shown an accompanying picture of the animal. Thus, the experiment contained three variables: context, animal, and picture.

After reading the vignette, participants answered several questions designed to assess their agreement with Rowe's judgment that the suffering described in the vignette was pointless. Since Rowe understands pointless suffering as suffering that does not bring about some greater good or prevent some equal or greater evil, the earlier study focused especially on items asking participants whether the situation in the story might bring about such a good or prevent such an evil. As we will discuss below, questions remain about how best to measure agreement with Rowe. Still, the earlier study provided an initial empirical basis for investigating how people respond to Rowe-style cases.

The results were striking. The animal described in the vignette and the presence or absence of a picture did not have a measurable effect on participants' responses. By contrast, the inclusion of broader background information significantly diminished perceptions of pointlessness: participants who received the additional background information were less likely to agree with Rowe than participants who received the original vignette alone. This finding is what we call the *context effect*.

This result is philosophically important. If judgments about the apparent pointlessness of Rowe-style suffering are highly sensitive to whether the suffering is presented in isolation or situated within broader background information, then this raises questions about how much evidential weight such judgments should carry. Perhaps the contextless presentation of the fawn case elicits an appearance of pointlessness that would be weakened, or even defeated, once the case is embedded in a wider frame. Given the exploratory nature of the earlier research, however, this conclusion must remain tentative. More importantly for present purposes, the earlier study left unanswered a further question: *why* does context have this effect? Does context diminish perceived pointlessness because it provides a possible point or explanation for the suffering? Because the original background information was positively valenced? Because it distracted participants from the suffering? Or maybe because it made salient the limits of our epistemic position? The remainder of this paper investigates these possibilities.

## §2: Questions and Hypotheses

Serious questions remain about the accuracy of these findings and what to make of them. To start, there is the question of whether Church, Warchol, and Barrett's (2022) two-item index is an accurate metric for agreement with Rowe. The challenge of making sure what is measured lines up with what philosophers are talking about is a perennial challenge for any work within experimental philosophy, and the situation is no different here.

For instance, participants might agree that some greater good *could* be accomplished (or some worse evil prevented) by the fawn's suffering but think of this possibility as extremely remote.<sup>5</sup> To disagree with Rowe, however, participants need to think of this possibility as being plausible, or at least not highly unlikely.

To address this concern, our team developed a nine-item index to measure agreement with Rowe, which we used in the present study. First, the two from Church et al. (2022):

---

<sup>5</sup> I am grateful to Stephen Wykstra for pointing out this objection.

“Some equal or greater evil could have been prevented because of the situation in the story.”  
“Some equal or greater good could be accomplished because of the situation in the story.”

Plus the following seven items:

“It is plausible that the situation in the story brings about some greater good.”  
“It is plausible that the situation in the story prevents some worse evil.”  
“Probably, the situation in the story doesn’t bring about some greater good.” (*reverse scored*)  
“Probably, the situation in the story doesn’t prevent some worse evil.” (*reverse scored*)  
“It is reasonable to believe that the situation in the story doesn’t bring about some greater good” (*reverse scored*)  
“It is reasonable to believe that the situation in the story doesn’t prevent some worse evil.” (*reverse scored*)  
The story you just read is an example of pointless suffering.

This expanded index allows us to better gauge how credible participants find the possibility of a greater good or worse evil. And since it includes the original 2-item index, it also allows them to see whether their original findings replicate.

But what explains the context effect identified by Church, Warchol, and Barrett (2022)? One possibility is that the context given to participants was too positive and that this positive valence is what diminishes perceptions of pointlessness. Call this the *positive valence hypothesis*. Or maybe the context that was given to participants contained an implicit theodicy. On this proposal, contexts drive down the perceptions of pointlessness only (or primarily) when they invoke a theodicy. Call this the *implicit theodicy hypothesis*. Or we might think that the context is simply distracting participants from the pointlessness of the suffering. So again, on this hypothesis, it is not the context *per se* that’s driving down the perception of pointlessness in Rowe’s vignette. That is, for example, it’s not as if participants are somehow seeing a “point” in the suffering. Instead, the sense of felt pointlessness is just being lost in the shuffle of new information. Call this the *distraction hypothesis*. And, finally, we might worry that the effect of context is merely a result of making participants aware of how little they know regarding the situation surrounding Rowe’s vignette and that it is this awareness, and not the context itself, that is driving down perceptions of pointlessness. Call this the *skeptical theism hypothesis*.<sup>6</sup>

In the present study, our team aimed to explore each of these hypotheses by introducing new vignette conditions. Let’s call the original two vignette types the “Regular Context Condition” and the “No Context Condition.” The following three conditions were also added to our study: “Negative Condition,” “Theodicy Condition,” “Junk Condition” to test the first three of those hypotheses above.

Addressing the positive valence hypothesis, the Negative Condition gave participants a negatively valenced vignette alongside the target case of suffering (i.e., absolutely nothing good is reported to come about from the fawn’s suffering, only bad). If the Negative Condition significantly diminishes the perception of pointlessness, then we’ll have reason to think that it’s not just the more positive valence of the Regular Context Condition that was causing participants to report diminished perceptions of pointlessness.

Addressing the implicit theodicy hypothesis, the Theodicy Condition presented a full-throated theodicy alongside the original vignette. If the implicit theodicy hypothesis were true, we

---

<sup>6</sup> For more on skeptical theism as a general response to the problem of evil, see Perrine 2023. For a skeptical-theist response to Rowe’s formulation of the problem of evil in particular, see Bergmann 2001.

would expect to see the Theodicy Condition significantly diminish the perception of pointlessness, even more than the Regular Context Condition. If the Theodicy Condition doesn't diminish perceptions of pointlessness as significantly as the Regular Context Condition, then we may doubt that the implicit theodicy hypothesis best explains the observed contextual effect.

Addressing the distraction hypothesis, the Junk Condition provided participants with contextual information that is only tangentially related to the target example of suffering (e.g., how forest fires produce light, that light travels nearly 300,000 km per second, etc.); if this condition significantly reduces the perception of pointlessness, then that would lend further credence to the distraction hypothesis. If the Junk Condition doesn't significantly diminish participants' perception of pointlessness, then it seems unlikely that the distraction hypothesis best explains the context effect.

Finally, in response to the skeptical theism hypothesis, some participants were randomly assigned to what we're calling the Skeptical Theism Condition. Participants who were assigned this condition read a brief discussion on the significant limitations humans face in trying to understand the suffering we encounter in this world before being asked to reflect on the original fawn vignette. If participants who were assigned the Skeptical Theism Condition report a significantly diminished perception of pointlessness in the target suffering, then that would lend credence to the skeptical theism hypothesis. If, however, subjects who were assigned to the Skeptical Theism Condition don't report a significantly diminished perception of pointlessness, then we'll have reason to think that the skeptical theism hypothesis does not best explain the observed contextual effect.

But maybe none of these hypotheses—*positive valence*, *implicit theodicy*, *distraction*, or *skeptical theism*—best explain the role context played in diminishing perceptions of pointlessness. This further research could suggest that the role context is playing is more mysterious than we might have previously expected.

### §3: Methods

This study was conducted online using Qualtrics to recruit participants and present the stories and all questions. Based upon this study's predecessor, to have sufficient statistical power, our recruiting target was 840 participants, randomly assigned to one of the aforementioned 6 story conditions: No Context, Regular Context Condition, Junk Condition, Negative Condition, Skeptical Theism Condition, and Theodicy Condition. The text of all six stories appears in the appendix below.

Our pre-registered<sup>7</sup> overarching hypothesis was that indicators of agreement with Rowe's judgment that the FAWN story is a case of pointless suffering will vary by condition type. More specifically, we hypothesized that scores on the dependent measure of agreement with Rowe (see below) would be in this order:

No Context > Junk > Negative > Skeptical Theism > Regular Context > Theodicy.

We also measured trait empathy and need for closure because we preregistered exploratory hypotheses concerning whether these individual-difference measures would predict agreement with Rowe. Because these measures are not central to the present paper's focus on vignette conditions, we report these analyses only briefly in §4.

---

<sup>7</sup> For purposes of anonymous review, the preregistration, study materials, and de-identified data are available via an anonymous view-only OSF link: [https://osf.io/v7nmh/overview?view\\_only=6fadd82a3c374e82813cd3f8fbbc6052](https://osf.io/v7nmh/overview?view_only=6fadd82a3c374e82813cd3f8fbbc6052). The public OSF record will be made available upon publication.

### *Participants*

A global sample of English-speaking participants were recruited using Qualtrics. The task was completed by 836 people, but 18 were eliminated for failing to attend to the task or taking too much or too little time to complete the survey.<sup>8</sup>

Of the 818 participants contributing to analyses, 513 claimed the United States as their nationality, 116 were Indian, and 189 were from numerous other nations.<sup>9</sup> The mean age was 39.28 years ( $SD=16.54$ , Median age = 36 years), and 491 reported being female, 314 male, and 13 reported other gender or not at all. Concerning education, 3.3% said they had completed 9th grade or less, 24.9 had a high school education, 20.2% claimed “some college or specialized training,” 11.6% reported an associates degree, 26.0% a bachelor’s degree, and 13.9% a master’s degree or higher. Education was not found to associate with any of the hypothesized relationships and will not be discussed further.

Of the 19 different categories of religions and worldviews provided as choices, only ten were selected by 2% or more of the participants: Spiritual but not Religious (18.6%), Jewish (15.3%), Hindu (13.6%), Other Christian (12.3%), Catholic (5.3%), Protestant (5.0%), Muslim (3.8%), Agnostic (2.7%), Atheist (2.4%), and Pagan (2.0%). “Other” was chosen by 13.9%. We had no hypotheses concerning religious or worldview identification and given the small size of each group across each story condition, we report no analyses using worldview as a variable.

### *Measures*

We asked participants questions to create a general profile of the sample (nationality, religious affiliation/worldview, educational attainment, gender). Time to complete the task was also measured to help identify people who did not give the task proper attention.<sup>10</sup>

A 17-item version of the Need for Closure Scale was used to measure the desire for conceptual certainty (Webster and Kruglanski 1994; Roets and Van Hiel 2007). Five items were reverse scored. Example items include, “Even after I’ve made up my mind about something, I am always eager to consider a different opinion” and “I dislike questions which could be answered in many different ways.” The 16-item Toronto Empathy Questionnaire was used to measure trait empathy (Spreng et al. 2009). Eight items were reverse scored. Examples include “When someone else is feeling excited, I tend to get excited too” and “I become irritated when someone cries.” For both scales, items were presented in a random order for each participant.

Nine items were used to index what Rowe meant by a case of “pointless suffering,” based upon our reading of Rowe. These items were placed on a 7-point agreement scale from -3 to +3:

1. Some equal or greater evil could have been prevented because of the situation in the story. (*reverse scored*)

---

<sup>8</sup> Six were eliminated for taking longer than three standard deviations more than the mean to complete. Five of these took more than four standard deviations longer than the mean. One was dropped for taking less than 5 minutes to complete, a speed judged too fast to adequately consider the items. Indeed, this participant answered the same for every item. All but one of the others omitted from analyses gave the same response for all 9 items of the Rowe index, all 16 items of the Toronto Empathy Questionnaire (TEQ), and all 17 items of the Need for Closure Scale (NECS). The final dropped participant similarly “flatlined” the TES, and the Rowe index, and answered 15 of the 17 TES as 5s, and 2 of the TES as 4s. As all three of these indexes have reverse-coded items, this pattern of responding appeared unconnected to the items and was more consistent with an attempt to complete the survey as quickly as possible, especially because the scores were at the extreme of the scales instead of the middle.

<sup>9</sup> Outside of India and the United States, no other nationality was well enough represented to constitute a meaningful subgroup for analytic purposes.

<sup>10</sup> After excluding extreme cases, time to completion was not significantly correlated with any of the main variables of interest.

2. Some equal or greater good could be accomplished because of the situation in the story. (*reverse scored*)
3. It is plausible that the situation in the story brings about some greater good. (*reverse scored*)
4. It is plausible that the situation in the story prevents some worse evil. (*reverse scored*)
5. Probably, the situation in the story doesn't bring about some greater good.
6. Probably, the situation in the story doesn't prevent some worse evil.
7. It is reasonable to believe that the situation in the story doesn't bring about some greater good.
8. It is reasonable to believe that the situation in the story doesn't prevent some worse evil.
9. The story you just read is an example of pointless suffering.

Much like listing the features of a great car (e.g., smooth and swift acceleration, great gas mileage, comfortable interior, ample cupholders, etc.), which need not cohere with each other, we made no assumption that these different dimensions of Rowe's pointlessness constitute anything like a coherent psychological construct. Nonetheless, to reduce the number of potential dependent variables, we calculated a Cronbach's alpha for these items (.61) and examined the correlation matrix. Items 1-4 all correlated with each other at  $r > .3$ , but with items 5-8  $r < .1$ , and vice versa. Reminiscent of the previous study by Church, Warchol, and Barrett (2022), the face-valid but ambiguous item 9 ("The story you read is an example of pointless suffering") only showed  $r > .3$  with item 5, and weak relationships with items 1-4. That is, the items phrased as supporting Rowe's claim to pointlessness grouped more strongly than those that require reverse coding. These clusterings motivated calculating Cronbach's alphas for items 1-4 (.68) and items 5-8 (.73). In keeping with our pre-registered hypotheses, we used the 9-item index of agreement with Rowe for our primary analyses. The 9 items were averaged after reverse-scoring items 1-4. Nonetheless, the stronger alphas of the two groups of items, one indicating general agreement with Rowe's judgment and one disagreeing with Rowe, motivated additional post-hoc analyses using their averages as dependent variables.

#### §4: Results

Although not central to the present paper's primary focus, we also tested whether trait empathy and need for closure predicted agreement with Rowe. Empathy did not significantly predict scores on the 9-item Rowe index. Need for Closure was a statistically significant but very weak predictor of the 9-item Rowe index, and exploratory analyses using the 4-item sub-indexes suggested that higher Need for Closure was associated with higher agreement with both pro- and anti-Rowe items. Given the small effect sizes and the ambiguity of these exploratory results, we do not interpret them further here.

The hypothesis that vignette conditions would impact Rowe agreement scores was tested by one-way ANOVA with the average of the nine Rowe items as the dependent variable,  $F(5,812) = 8.54, p < .001, \eta^2 = .05$ . Table 1 displays the counts, means, standard deviations, and one-sample t-tests against a value of zero for each of the six conditions. As values ranged from -3 to +3, a score of zero represented no clear decision to agree or disagree with the items. Positive scores indicate a tendency to agree with Rowe's characterization of the events of the story as an example of pointless suffering, unlikely to bring about a greater good or prevent a greater evil. As Table 1 illustrates, only the No Context Condition, Rowe's original FAWN story, had scores greater than zero. The Regular Context Condition, Skeptical Theism, and Theodicy conditions all had scores significantly less than

zero, suggesting a tendency to reject Rowe’s judgment. Figure 1 depicts the mean Rowe 9-item index scores across the six vignette conditions.

**Figure 1: Mean Rowe 9-item index score by vignette condition.** Error bars represent 95% confidence intervals. Scores above 0 indicate greater agreement with Rowe’s judgment that the fawn’s suffering is pointless; scores below 0 indicate greater disagreement with Rowe’s judgment.

[Insert Figure 1]

**Table 1: Rowe 9-item Index ANOVA by Condition, One-sample t-tests against 0.**

Condition	Count	Mean	SD	t	df	p
No Context	140	.28	.83	3.97	139	<.001
Junk Condition	135	.02	.68	0.30	134	.764
Negative Condition	137	-.01	.69	0.19	134	.853
Theodicy	139	-.12	.62	2.21	138	.029
Skeptical Theism	134	-.13	.65	2.38	133	.019
Regular Context	133	-.21	.83	3.57	132	<.001

ANOVA:  $F(5,812) = 8.54, p < .001, \eta^2 = .05$

Planned contrasts detected that Rowe scores in the No Context Condition were higher than the next highest condition, Junk Condition,  $t(260) = 2.95, p = .002$ , and therefore also significantly greater than all other conditions. The Regular Context Condition had significantly lower scores than the Junk Condition,  $t(265) = 2.79, p = .006$ ; and the Negative Condition,  $t(266) = 2.35, p = .02$ ; but not than Theodicy or Skeptical Theism, which also did not significantly differ from each other. Thus, the fine-grained prediction that No Context > Junk Context > Negative Context > Skeptical Theism > Regular Context > Theodicy, was not supported.

To further explore these results, two more one-way ANOVAs were conducted with accompanying one-sample t-tests. One ANOVA used the four-item index of agreement with Rowe’s judgment as the dependent variable, and one used the four-item index of disagreement. Summaries of these tests appear in Tables 2 and 3. These exploratory analyses were not pre-registered.

**Table 2: Rowe Agreement 4-item Index ANOVA by Condition, One-sample t-tests against 0.**

Condition	Count	Mean	SD	t	df	p
No Context	140	.59	1.12	6.22	139	<.001
Junk Condition	135	.36	.96	4.38	134	<.001
Negative	137	.35	.92	4.44	136	<.001
Theodicy	139	.26	1.03	2.96	138	.004
Regular Context	133	.22	.98	2.64	132	.009
Skeptical Theism	134	.14	.86	1.90	133	.059

ANOVA:  $F(5,812) = 3.40, p = .005, \eta^2 = .02$

**Table 3: Rowe Disagreement 4-item Index ANOVA by Condition, One-sample t-tests against 0.**

Condition	Count	Mean	SD	t	df	p
No Context	140	.03	1.18	0.29	139	.775

Junk Condition	135	.35	1.07	3.85	134	<.001
Negative	137	.39	1.10	4.17	136	<.001
Skeptical Theism	134	.40	.99	4.75	133	<.001
Theodicy	139	.51	1.01	6.00	138	<.001
Regular Context	133	.66	.93	8.17	132	<.001

ANOVA:  $F(5,812) = 5.51, p < .001, \eta^2 = .03$

Both ANOVAs detected significant differences among conditions. Participants answered significantly above zero, indicating agreement with the items, in five of six conditions. For the index expressing a tendency to agree that the situation in the story was not bringing about a greater good or preventing a greater evil (i.e., was ‘pointless suffering’), only in the Skeptical Theism condition were scores not greater than zero. For the index expressing disagreement with Rowe’s judgment—that is, agreement that the situation in the story may have brought about a greater good or prevented a greater evil—all but the No Context Condition were significantly greater than zero. That is, whether the items expressed agreement or disagreement with Rowe, participants in most conditions were inclined to agree with the items presented to them.

*Replication of McAllister et al. (2024).*

In McAllister et al. (2024), two items from the 4-item “disagreement” index constituted the dependent variable. To check for replication here, the same two conditions, No Context and Regular Context, were compared with the same dependent variable, the average of “Some equal or greater evil could have been prevented because of the situation in the story,” and “Some equal or greater good could be accomplished because of the situation in the story.” As a variable, positive scores indicate seeing some potential good arising from the FAWN story, and, thus, represent disagreement with Rowe’s judgment that it is a case of pointless suffering. The independent samples t-test detected a difference between conditions as previously,  $t(265) = 3.68, p < .001, r^2 = .014$ . The Regular Context Condition had higher scores ( $M = .68, SD = 1.07$ ) than the No Context Condition ( $M = .15, SD = 1.31$ ). That is, the additional context in the Regular Context condition appeared to elevate the confidence that some good could come from the event in the story.

*Results Summary.*

The vignette conditions were found to impact agreement or disagreement with Rowe’s judgment regarding the fawn in the forest. This study’s predecessor, using just two items and two context conditions, was also replicated. We observe, however, that all of these significant effects are small in size. For this reason, the most philosophically important result may be the broader descriptive pattern: across the full sample, no measure of agreement with Rowe consistently scored significantly above the scale midpoint, which indicates neither agreement nor disagreement. We discuss the significance of this pattern below.

## §5: Discussion

The present study reinforces a central finding from recent experimental work on Rowe-style cases: judgments about whether a target instance of suffering is pointless are sensitive to how the case is presented. More specifically, we found that the vignette condition significantly shaped participants’ agreement with Rowe’s judgment that the fawn’s suffering was pointless. We also replicated the earlier finding that the Regular Context Condition reduces agreement with Rowe when compared to the No Context Condition. This replication is important. The earlier results suggested that adding

broader background information to Rowe's fawn vignette diminished perceptions of pointlessness. The present study provides further evidence that this effect is not merely an artifact of the earlier study's design or measurement strategy.

At the same time, the present results also counsel restraint. The observed effects were small, and the descriptive statistics remain philosophically interesting in their own right. Across conditions, participants did not consistently or strongly endorse Rowe's judgment above the midpoint of the scale. Indeed, the No Context Condition was the only condition in which the 9-item index rose significantly above the neutral midpoint. In the Regular Context Condition, Skeptical Theism Condition, and Theodicy Condition, participants' responses fell significantly below the midpoint, suggesting a tendency to reject Rowe's judgment. The broader picture, then, is not that participants strongly agreed with Rowe except when context intervened. Rather, participants were often ambivalent or inclined to disagree with Rowe, and the No Context Condition was the condition most favorable to Rowe's assessment of the fawn's suffering.

This matters for how the present findings should be interpreted. The central result is not simply that context "solves" the problem posed by Rowe's fawn. Nor is it that broader background information shows that the fawn's suffering really is justified. The result is more modest but still significant: the appearance of pointlessness in Rowe-style cases appears to be presentation-sensitive. The very same target suffering can elicit different responses depending on whether it is presented in isolation or situated within broader background information. Since Rowe's argument depends, at least in part, on the apparent pointlessness of cases like the fawn's suffering, this presentation-sensitivity raises important questions about what such appearances are tracking and how much evidential weight they should bear.

### *5.1. Evaluating the Hypotheses*

The present study was designed to test several hypotheses about why the Regular Context Condition diminished perceived pointlessness. One possibility was that the effect was driven by the positive valence of the original background information. In the earlier study, the added background concerned the role that smaller forest fires can play in forest ecosystems: clearing dead organic material, enriching soil, facilitating new growth, and preventing more catastrophic fires. This information is not a theodicy, at least not explicitly, but it is easy to see why one might regard it as optimistic or positively valenced. Perhaps, then, participants were less inclined to agree with Rowe simply because the background information framed forest fires in a more positive light.

The results do not support that explanation. The Negative Condition did not increase agreement with Rowe relative to the No Context Condition in the way one would expect if positive valence were the key driver of the original effect. Indeed, the Negative Condition did not significantly rise above the midpoint on the 9-item index. This suggests that the context effect is not simply a function of optimistic or positive framing. Something more than positive valence appears to be involved.

A second possibility was that the Regular Context Condition diminished perceived pointlessness because it contained an implicit theodicy. On this hypothesis, participants were not responding to context as such; rather, they were responding to the suggestion that the suffering might serve some greater ecological purpose. If this were correct, then one might expect an explicit theodicy to produce an even stronger reduction in agreement with Rowe. After all, if the relevant mechanism is the provision of a point, then a condition that provides a full-throated point should have an especially strong effect.

Here too, however, the results complicate the hypothesis. The Theodicy Condition did reduce agreement with Rowe relative to the No Context Condition, but it did not clearly outperform the Regular Context Condition. This is important. The Regular Context Condition did not explicitly

present itself as a theodicy, did not ask participants to think about God (in a sample where a majority of subjects were theistic), and did not directly argue that the fawn's suffering was justified. Yet it produced an effect at least comparable to the Theodicy Condition. This gives us reason to doubt that the original context effect is best explained by the presence of an implicit theodicy. Theodical reflection may matter, but it does not appear to be the unique or primary explanation of the effect.

A third possibility was that the Regular Context Condition reduced agreement with Rowe by distracting participants from the suffering. On this hypothesis, additional information does not diminish the perception of pointlessness because participants are finding a point in the suffering, or even because they are becoming more cautious about judging it pointless. Rather, the added information simply shifts attention away from the horrific details of the fawn's agony. The felt force of the case is thereby weakened or "lost in the shuffle."

The Junk Condition was designed to test this possibility. Participants in this condition received additional information related to fire, but not in a way that was meaningfully connected to the suffering described in Rowe's vignette. If distraction were the main mechanism, then the Junk Condition should have produced an effect similar to the Regular Context Condition. But it did not. Although the Junk Condition differed from the No Context Condition, it did not reproduce the effect of the Regular Context Condition: unlike the Regular Context Condition, the Junk Condition did not fall significantly below the neutral midpoint. This suggests that the original context effect is not merely a distraction effect. The kind of additional information matters. Mere informational bulk does not seem sufficient.

Finally, we considered a skeptical-theism hypothesis. Perhaps context diminishes perceived pointlessness because it makes salient the limits of our epistemic position. On this view, participants are not finding a point in the suffering. Instead, they are being reminded that they do not know enough to confidently judge that there is no point. If so, then a condition explicitly prompting reflection on human cognitive limitation should reproduce the effect of the Regular Context Condition.

The Skeptical Theism Condition did reduce agreement with Rowe relative to the No Context Condition. This result is philosophically interesting, and it suggests that epistemic-limitation considerations may have some effect on judgments about Rowe-style cases. Still, skeptical-theist prompting does not fully explain the original context effect. The Skeptical Theism Condition did not clearly outperform the Regular Context Condition, and it did not uniquely account for the overall pattern of results. Reflection on epistemic limitation may be part of the story, but it does not appear to be the whole story.

Taken together, these results suggest that none of the most natural explanations fully accounts for the context effect. The effect is not reducible to positive valence, explicit theodicy, distraction, or skeptical-theist prompting. This is not a disappointing result. On the contrary, it is one of the most philosophically interesting results of the study. It suggests that context matters, but that its role is more subtle than we initially expected.

### *5.2. Context, Intelligibility, and Pointlessness*

What, then, is context doing? The present study does not allow us to answer this question definitively. Still, the pattern of results suggests one promising possibility: context may diminish perceived pointlessness by rendering the target suffering less isolated and more intelligible.<sup>11</sup> That is,

---

<sup>11</sup> This suggestion is broadly consonant with psychological work on meaning-making in response to stressful life events. Park and Folkman (1997), for example, distinguish between global meaning and situational meaning and emphasize the

context may matter not because it explicitly justifies the suffering but because it situates the suffering within a broader frame of causes, consequences, and possible significance.

This distinction is important. A background frame need not show that suffering is justified in order to weaken the judgment that it is pointless.<sup>12</sup> In Rowe's original vignette, the fawn's suffering is presented in stark isolation. The case is brief, vivid, and devastating. A fawn is trapped in a fire, suffers horribly for several days, and dies. Nothing in the vignette indicates what preceded the fire, what follows from it, what role the fire plays in the surrounding ecosystem, or how the event fits into a larger causal order. The suffering appears as a nearly self-contained episode of agony. In that presentation, the judgment of pointlessness can seem almost immediate.

The Regular Context Condition changes this presentation. It does not remove the suffering. It does not deny the horror of the fawn's agony. Nor does it show that the fawn's suffering was required for some greater good. But it does alter the frame in which the suffering is encountered. The suffering is no longer presented as an isolated episode. It is embedded within a broader account of forest fires, ecological renewal, causal processes, and long-term consequences. This broader frame may weaken the sense that one can read pointlessness directly off the event itself.

This helps explain why the Junk Condition did not have the same effect. The Junk Condition adds material, but it does not render the target suffering more intelligible in the relevant way. It supplies facts, but not a meaningful frame. It is therefore not surprising that mere additional information failed to reproduce the effect of the Regular Context Condition. What matters may not be quantity of information but the kind of intelligibility that the information provides.

The same point may also explain why the Negative Condition did not simply reverse the effect. Even negatively valenced context can situate an event within a broader frame. It may not make the event look good, and it may not provide any redemptive interpretation. But it can still make the event appear less isolated. The suffering becomes part of a broader causal story, even if that story is bleak. This suggests that context may influence judgments of pointlessness partly by altering the perceived completeness, causal connectedness, or intelligibility of the case, not merely by shifting its emotional valence.

This is a useful way of understanding the present results: context need not provide a justification in order to undermine confidence in pointlessness. It may be enough for context to make the target event appear less epistemically self-contained. When the fawn's suffering is presented without broader background, participants may feel invited to judge the event as a whole from the brief description alone. When broader background is supplied, that invitation is weakened. Participants may become less confident that the information provided is sufficient to support Rowe's conclusion.

### *5.3. Narrative as a Promising Explanatory Framework*

One natural way to develop this thought is through the lens of narrative. The present study did not directly test a narrative hypothesis. We did not manipulate narrative structure as such, nor did we measure participants' judgments of narrative coherence, causal integration, or perceived completeness. Still, the results are consistent with the idea that human judgments about suffering are

---

role of reappraisal in bringing the appraised meaning of an event into relation with broader beliefs and goals. Park (2010) further distinguishes between meaning-making as a search for comprehensibility and as a search for significance.

<sup>12</sup> One can come to understand an event as less isolated, less opaque, or more integrated into a broader causal or interpretive frame without thereby judging the event good, justified, or redemptive. This distinction also leaves room for more sufferer-centered accounts of the defeat or redemption of evil. Adams (1999), for example, argues that horrendous evils threaten to defeat the positive meaning of a participant's life, and that any adequate divine response must defeat such evils in relation to the participant's own life, rather than merely by appeal to goods considered abstractly.

shaped by whether suffering is presented as an isolated episode or embedded within a broader narrative or quasi-narrative frame.

This suggestion fits with a wide range of work in philosophy and psychology. Jerome Bruner famously distinguishes between paradigmatic and narrative modes of thought (Bruner 1986), later developing this account by arguing that narrative is a basic way human beings organize experience and construct reality (Bruner 1991). On this view, human beings do not understand the world merely by subsuming events under abstract laws or principles. We also understand events by placing them into stories: temporally extended, causally structured, meaning-laden accounts of agents, actions, consequences, and significance. Narrative is not merely decorative. It is one of the basic forms through which human beings render experience intelligible.

Psychological work on narrative identity points in a similar direction. Dan McAdams, Kate McLean, and others have argued that human beings make sense of their lives by constructing “internalized and evolving life [stories]” that confer unity, purpose, and significance (McAdams and McLean 2013, 233). In this literature, suffering and adversity are often understood not merely as discrete negative events, but as episodes that can be integrated into broader structures of meaning. This does not imply that all suffering is good, justified, or redemptive. But it does suggest that human beings often process suffering differently when it is located within a larger temporal and causal frame.

Eleonore Stump’s work on narrative and suffering provides an especially relevant philosophical analogue. In *Wandering in Darkness* (2010), Stump argues that the problem of suffering cannot be adequately understood through abstract propositions alone. Narratives, on her view, can disclose morally and spiritually significant features of suffering that are not easily captured by thin, context-stripped descriptions. One need not accept Stump’s broader theodicy to appreciate the methodological point: judgments about suffering may depend on details of context, character, history, relationship, and narrative structure that compressed philosophical examples often omit.

This point is directly relevant to Rowe’s fawn case. Rowe’s vignette is deliberately sparse. Its sparseness is part of its philosophical power. But that same sparseness may also be part of what generates the appearance of pointlessness. The case presents suffering in a form that is almost maximally decontextualized. It strips away the very sorts of background that, in ordinary life, often shape our judgments about whether an event has a point, role, consequence, or meaning. If human beings ordinarily assess suffering by embedding it in broader structures of intelligibility, then the contextless fawn vignette may elicit a judgment under conditions that are psychologically and philosophically unusual.

This is not to say that Rowe has made a methodological mistake. Philosophical thought experiments often work by isolating a target phenomenon. The point of a vignette is frequently to strip away irrelevant detail so that the philosophically important structure can be seen more clearly. But the present findings suggest that, in cases of suffering and pointlessness, it is not obvious that the stripped-away material is irrelevant. If judgments of pointlessness depend partly on the presence or absence of broader causal, temporal, and narrative frames, then context is not merely ornamental. It may be part of what makes the judgment possible.

#### *5.4. Implications for Rowe’s Argument*

What do these findings imply for Rowe’s evidential argument from evil? The answer must be stated carefully. The present study does not claim to show that Rowe’s argument fails. It does not show that the fawn’s suffering is not pointless. Nor does it show that there is some greater good or worse evil that justifies God in permitting such suffering. The data are psychological, not metaphysical. They tell us something about how people judge cases of suffering, not directly about whether those judgments are true.

Even so, the findings have philosophical significance. Rowe's argument depends on the claim that certain instances of suffering appear pointless, and that this appearance provides some evidence that they really are pointless. If the appearance of pointlessness is highly sensitive to features of presentation, then we have reason to ask how evidentially stable that appearance is. A judgment that arises strongly in one presentation but weakens substantially in another may still be true, but its evidential status becomes more complicated.

This is especially important because Rowe's original vignette is not merely one presentation among many. It is a contextless presentation. The fawn's suffering is presented without any broader account of the causal or temporal structure in which the event occurs. If that kind of presentation is especially likely to elicit judgments of pointlessness, then philosophers should be cautious about moving too quickly from the felt force of the vignette to stronger conclusions about the existence of gratuitous evil.

Of course, a defender of Rowe might respond that this is exactly as it should be. Perhaps adding context diminishes perceived pointlessness only because human beings are prone to overinterpretation. Perhaps people are inclined to find patterns, purposes, or possible goods where none exist.<sup>13</sup> On this response, the contextless vignette might be epistemically superior precisely because it prevents participants from inventing a point. Broader context, on this view, may obscure rather than clarify.

That response remains available, and the present study does not rule it out. But the results do make the dialectical situation more complex. The defender of Rowe cannot simply assume that the appearance of pointlessness generated by the sparse vignette is the privileged one. If context changes the judgment, then we need some reason to think that the judgment elicited by the contextless presentation is more reliable than the judgment elicited by the contextualized presentation. Without such a reason, the evidential force of the fawn case is less straightforward than it may initially appear.

The central lesson, then, is not that context defeats Rowe. It is that context complicates the evidential role of Rowe-style intuitions. The appearance of pointlessness may be real, powerful, and philosophically important. But it may also be sensitive to the way suffering is framed. That sensitivity should make us more cautious in treating such appearances as stable evidence for the existence of gratuitous suffering.

### *5.5. Implications for Experimental Philosophy*

The present findings also have broader implications for experimental philosophy. Experimental philosophers often use brief vignettes to elicit judgments about philosophically significant cases. This method has many virtues. Brief vignettes are efficient, controllable, and capable of isolating variables that might otherwise be confounded. But the present study highlights a limitation: some judgments may depend on precisely the kind of background that brief vignettes tend to omit.

This worry is not unique to philosophy of religion. Judgments about moral responsibility, personal identity, knowledge, free will, and harm may all be sensitive to how much contextual and narrative information is provided.<sup>14</sup> If a judgment concerns whether an event is meaningful, justified, excusable, identity-preserving, or pointless, then it may be especially dependent on broader causal

---

<sup>13</sup> See, for example, Kelemen 1999; Shermer 2008.

<sup>14</sup> For a related methodological point, see Schechtman 2022. Schechtman argues that the brevity and clarity of experimental vignettes is one of their strengths but that this same feature also imposes limitations: richer fictional narratives can explore "forms of messiness and ambiguity" as well as individual details that are difficult to capture experimentally (2022, 99). Her discussion focuses especially on personal identity and morality, but the broader methodological point is relevant here: some philosophical judgments may be sensitive to narrative and contextual features that sparse vignettes necessarily omit.

and narrative structure. In such cases, brevity may not merely simplify the case; it may alter the phenomenon being studied.

This suggests a methodological challenge. Experimental philosophy often seeks to determine what people think about a philosophically important case. But there may be no single neutral way to present the case. A sparse presentation may seem methodologically clean, but it may also strip away information that participants would ordinarily use to form the relevant judgment. A richly contextualized presentation may seem more ecologically valid, but it may introduce new interpretive cues that shape the response. Neither approach is obviously privileged in advance.

The way forward is not to abandon vignettes, but to become more precise about what different vignette forms are testing.<sup>15</sup> A contextless vignette may test how people respond to suffering when it is presented in isolation. A background condition may test how people respond when the same suffering is placed within a broader causal frame. A narrative condition may test how judgments change when the suffering is embedded within a temporally extended story. These are different psychological and philosophical tasks. Treating them as such would allow future work to investigate not simply whether people agree with Rowe, but what features of presentation shape that agreement.

### *5.6. Future Directions*

The present study therefore opens several avenues for future research. Most importantly, future studies should test the narrative hypothesis directly. Rather than contrasting context with no context in a broad way, subsequent work should manipulate specific dimensions of narrative embedding. For example, researchers might vary whether the same suffering is presented as an isolated event, as part of a causally connected sequence, or as part of a temporally extended story with before-and-after structure. If narrative embedding is doing some of the work, then causally and temporally integrated presentations should reduce perceived pointlessness more reliably than equally long but disconnected background information.

Future work should also distinguish between redemptive and non-redemptive forms of narrative. One might suspect that suffering appears less pointless only when it is embedded within a redemptive story, where the suffering leads to some good outcome. But the present study suggests that this may be too simple. Even the Negative Condition, where the target suffering explicitly did not seem to lead to any greater good or prevent any greater evil, did not straightforwardly increase agreement with Rowe. This raises the possibility that intelligibility and redemption can come apart. A story may render an event intelligible without rendering it good. Future studies should therefore compare redemptive, tragic, unresolved, and merely causal forms of narrative embedding.

It would also be useful to measure more than agreement with Rowe. The present study already suggests that “pointlessness” is not psychologically simple. Participants may be making any among several distinct judgments: whether the suffering actually has no point, whether it appears to have no point, whether we are in a position to know that it has no point, whether the event is causally intelligible, whether the case feels complete, and whether the suffering is emotionally overwhelming. Future work should measure these dimensions separately. Doing so would help determine whether context reduces perceived pointlessness by increasing causal intelligibility, reducing epistemic confidence, lowering affective intensity, or some combination of these.

Finally, future research should explore whether the context effect generalizes beyond Rowe’s fawn case. The fawn vignette is especially important because of its role in the contemporary philosophy of religion literature, but the underlying phenomenon may be much broader. Judgments

---

<sup>15</sup> Additionally, it will be worth considering the cognitive genesis of the intuitions these vignettes elicit. See De Cruz 2015.

about apparently pointless suffering arise in many domains: animal suffering, natural disasters, illness, bereavement, injustice, and horrendous evil. If context and narrative shape perceptions of pointlessness across these domains, then the implications of the present research extend beyond Rowe's argument. They may illuminate something general about how human beings interpret suffering.

## Conclusion

The central conclusion of this study is modest but important. Context matters. More precisely, the way suffering is presented affects whether people judge that suffering to be pointless. The present study replicates earlier evidence for this context effect while showing that none of the most natural explanations fully accounts for it. The effect does not appear to be reducible to positive valence, explicit theodicy, distraction, or skeptical-theist reflection.

This leaves us with a puzzle. Why does context diminish perceived pointlessness? The present study does not settle the question, but it suggests that the answer may have to do with intelligibility. Context may matter because it prevents suffering from appearing as a radically isolated event. It embeds suffering within a broader frame of causes, consequences, and possible significance. Narrative is one especially important form such framing can take.

If this is right, then the philosophical significance of the context effect is considerable. Rowe's fawn case derives much of its force from the apparent pointlessness of the suffering it describes. But if that appearance is sensitive to whether the suffering is presented in isolation or embedded within a broader frame, then the evidential role of the case is less straightforward than it first appears. This does not refute Rowe's argument—or at least that's not the claim of this paper—but it does suggest that philosophers should be cautious about treating the apparent pointlessness of sparse vignettes as a stable datum. The appearance of pointlessness may tell us something important about suffering. But it may also tell us something important about the human need to understand suffering in context.

**AI Use Declaration:** The authors used ChatGPT (OpenAI; GPT-5.5 Thinking) for limited assistance with manuscript preparation, including brainstorming, prose revision, copyediting, and clarification of wording. The tool was not used to collect data, conduct statistical analyses, generate research findings, or determine the paper's conclusions. All AI-assisted text was reviewed, revised, and approved by the authors, who take full responsibility for the content of the manuscript.

## Bibliography

- Adams, Marilyn McCord. 1999. *Horrendous Evils and the Goodness of God*. Cornell University Press.
- Bergmann, Michael. 2001. "Skeptical Theism and Rowe's New Evidential Argument From Evil." *Noûs* 35 (2): 278–96. <https://doi.org/10.1111/0029-4624.00297>.
- Bruner, Jerome. 1986. *Actual Minds, Possible Worlds*. Harvard University Press. <https://www.jstor.org/stable/j.ctv26070s7>.

- Bruner, Jerome. 1991. "The Narrative Construction of Reality." *Critical Inquiry* 18 (1): 1–21. <https://doi.org/10.1086/448619>.
- Church, Ian M., Rebecca Carlson, and Justin Barrett. 2020. "Evil Intuitions? The Problem of Evil, Experimental Philosophy, and the Need for Psychological Research." *Journal of Psychology and Theology*.
- Church, Ian M., Isaac Warchol, and Justin L. Barrett. 2022. "The Context of Suffering: Empirical Insights into the Problem of Evil." *TheoLogica: An International Journal for Philosophy of Religion and Philosophical Theology*.
- De Cruz, Helen. 2015. "Where Philosophical Intuitions Come From." *Australasian Journal of Philosophy* 93 (2): 233–49.
- Kelemen, D. 1999. "The Scope of Teleological Thinking in Preschool Children." *Cognition* 70 (3): 241–72. [https://doi.org/10.1016/s0010-0277\(99\)00010-4](https://doi.org/10.1016/s0010-0277(99)00010-4).
- McAdams, Dan P., and Kate C. McLean. 2013. "Narrative Identity." *Current Directions in Psychological Science* 22 (3): 233–38. <https://doi.org/10.1177/0963721413475622>.
- McAllister, Blake, Ian M. Church, Paul Rezkalla, and Long Nguyen. 2024. "Empirical Challenges to the Evidential Problem of Evil." In *The Oxford Studies in Experimental Philosophy*, edited by Joshua Knobe and Shaun Nichols. Oxford, UK: Oxford University Press.
- Park, Crystal L. 2010. "Making Sense of the Meaning Literature: An Integrative Review of Meaning Making and Its Effects on Adjustment to Stressful Life Events." *Psychological Bulletin* 136 (2): 257–301. <https://doi.org/10.1037/a0018301>.
- Park, Crystal L., and Susan Folkman. 1997. "Meaning in the Context of Stress and Coping." *Review of General Psychology* 1: 115–44. <https://doi.org/10.1037/1089-2680.1.2.115>.
- Perrine, Timothy. 2023. "Skeptical Theism." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman. Stanford University. <https://plato.stanford.edu/archives/spr2023/entries/skeptical-theism/>.
- Plantinga, Alvin. 2000. *Warranted Christian Belief*. Oxford University Press.
- Roets, Arne, and Alain Van Hiel. 2007. "Separating Ability From Need: Clarifying the Dimensional Structure of the Need for Closure Scale." *Personality and Social Psychology Bulletin* 33 (2): 266–80. <https://doi.org/10.1177/0146167206294744>.
- Rowe, William L. 1979. "The Problem of Evil and Some Varieties of Atheism." *American Philosophical Quarterly* 16 (4): 335–41.
- Rowe, William L. 1996. "The Evidential Argument From Evil: A Second Look." In *The Evidential Argument From Evil*, edited by Daniel Howard-Snyder. Indiana University Press.

- Schechtman, Marya. 2022. "The Whole Story: Identity and Narrative." In *Experimental Philosophy of Identity and the Self*, edited by Kevin Tobia. Bloomsbury.  
<https://philarchive.org/rec/SCHTWS-5>.
- Shermer, Michael. 2008. "Patternicity: Finding Meaningful Patterns in Meaningless Noise." *Scientific American*, December 1. <https://www.scientificamerican.com/article/patternicity-finding-meaningful-patterns/>.
- Spreng, R. Nathan, Margaret C. McKinnon, Raymond A. Mar, and Brian Levine. 2009. "The Toronto Empathy Questionnaire." *Journal of Personality Assessment* 91 (1): 62–71.  
<https://doi.org/10.1080/00223890802484381>.
- Stump, Eleonore. 2010. *Wandering in Darkness: Narrative and the Problem of Suffering*. Oxford University Press.
- Webster, D. M., and A. W. Kruglanski. 1994. "Individual Differences in Need for Cognitive Closure." *Journal of Personality and Social Psychology* 67 (6): 1049–62.  
<https://doi.org/10.1037//0022-3514.67.6.1049>.
- Wykstra, Stephen J. 1996. "Rowe's Noseeum Arguments From Evil." In *The Evidential Argument From Evil*, edited by Daniel Howard-Snyder. Indiana University Press.

## Appendix: The Six Context Types

*No Context:* Suppose in some distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering.

*Regular Context:* Forest fires are often viewed as some of the most dangerous and destructive natural disasters. While some fires of catastrophic size can be detrimental to forests and endanger human lives and infrastructure, smaller forest fires are actually an essential aspect of the forest ecosystem. It may seem counterintuitive that fires could be beneficial to the life of a forest, however, recent ecological research has shown that small burns play a major role in the health of an ecosystem as a whole. Fires, often resulting from lightning strikes, quickly and efficiently clear away thick undergrowth, dying trees, and the dead material that congregates on the forest floor. If left unchecked, dead organic material and undergrowth will prevent new trees and plants from taking root and being able to grow. The burnt organic material such as plants, shrubs, and animals, leave behind topsoil that is rich in nutrients from which new plant life can easily grow. Small forest fires also play an important role in preventing fires from reaching catastrophic sizes. When a fire is small, it is usually confined to burning the undergrowth and dead material on the forest floor and does not burn the tree canopy or kill the large trees of the forest. However, if a forest goes too long without a fire, the undergrowth will become so thick that when it does burn it will easily ignite not only the forest floor but also the trees themselves. Many experts attribute the record-setting fires that have been seen in recent years to decades of fire suppression in forests, which has left entire ecosystems

vulnerable to catastrophic fires. Many species of plants have adapted to occasional fires and can quickly regrow burnt branches. Some trees even need fire to reproduce due to seed-cones that will only open when exposed to extreme temperature.

Now, please consider the following story: suppose in a distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering.

*Junk Condition:* Fires start when a flammable or a combustible material, in combination with a sufficient quantity of an oxidizer such as oxygen gas or another oxygen-rich compound (though non-oxygen oxidizers exist), is exposed to a source of heat or ambient temperature above the flash point for the fuel/oxidizer mix and is able to sustain a rate of rapid oxidation that produces a chain reaction. This is commonly called the fire tetrahedron. Fire cannot exist without all of these elements in place and in the right proportions. For example, a flammable liquid will start burning only if the fuel and oxygen are in the right proportions. Some fuel-oxygen mixes may require a catalyst, a substance that is not consumed, when added, in any chemical reaction during combustion, but which enables the reactants to combust more readily. Once ignited, a chain reaction must take place whereby fires can sustain their own heat by the further release of heat energy in the process of combustion and may propagate, provided there is a continuous supply of an oxidizer and fuel. If the oxidizer is oxygen from the surrounding air, the presence of a force of gravity, or of some similar force caused by acceleration, is necessary to produce convection, which removes combustion products and brings a supply of oxygen to the fire. Without gravity, a fire rapidly surrounds itself with its own combustion products and non-oxidizing gases from the air, which exclude oxygen and extinguish the fire. Because of this, the risk of fire in a spacecraft is small when it is coasting in inertial flight. This does not apply if oxygen is supplied to the fire by some process other than thermal convection. Fire can be extinguished by removing any one of the elements of the fire tetrahedron.

Now, please consider the following story: suppose in a distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering.

*Negative Condition:* Forest fires are becoming more frequent, widespread, and destructive as a result of human influences on the environment. Fires are now causing an additional 3 million hectares (7.4 million acres) of tree cover loss per year than they did in 2001, and many of these fires burn all or most of a forest's living overstory trees. Fires are not naturally occurring in tropical rainforests, but in recent years, as deforestation and climate change have degraded and dried out intact forests, fires have been escaping into standing tropical rainforests as well. Many corporations and governments are quick to take advantage of the newly cleared land to reduce their overhead when taking on new building projects (regardless of whether those building projects are needed). In the event of a major forest fire, a company or a local government might use the newly cleared land to justify a new building project—flattening and paving the forest to make way for another store in a popular franchise or maybe another shopping mall. Obviously, once paved no plants would be able to grow, which would further limit (even in a very small way) the earth's ability to process carbon, which worsens (again, even if in a small way) global warming. Adding to the tragedy of the situation, given rising inflation and growing economic woes, most new-builds are expected to go out of business within 10 years, sitting empty for years to come. Paving the forests, and taking on doomed building

projects only exacerbates environmental and economic woes—wasting limited resources, preventing regrowth, and generally making the world an uglier and more uninhabitable place.

Now, please consider the following story: suppose in a distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering.

*Theodicy Condition:* Forest fires and other natural disasters cause a lot of suffering. Many lives are destroyed each year. If there's a good God, then wouldn't God prevent these evil things from happening? Maybe. Or maybe the world is better off if some of these things are allowed to happen. To prevent all natural evils, God would have to violate the laws of nature over and over again. What makes fire useful also makes it dangerous. There's no way to avoid that danger unless God constantly intervenes to prevent you (and everything else) from being harmed. If God is constantly intervening to prevent fires, hurricanes, and other natural disasters, then there basically wouldn't be laws of nature anymore. The beauty and elegance of the natural order would be largely destroyed. We couldn't do science either, nor would there be any need to since our condition would be the same regardless. For instance, we wouldn't need to strive to grow better crops because God would never let there be too little food. In fact, we wouldn't even need to plant crops because God would just miraculously make food appear! No matter what we do, God won't let us face the natural consequences of our decisions. This leads to the greatest cost of eliminating natural evils, which is that it would render many of our choices meaningless. It doesn't matter whether we try to eliminate disease or take care of the planet, because, whatever we do, the outcome is the same. We'd also lose the opportunity to grow in the face of challenges. How can we be courageous if there is nothing to fear? How can we learn the importance of protecting others when there are no threats? How can we show compassion when no one is ever needy? If God intervened to prevent all natural disasters, then that would remove some of the very best things about our world. None of this assumes that there is a God, of course. My point is just that, if a good God does exist, I can understand why such a God might allow natural disasters to occur.

Now, please consider the following story: suppose in a distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering.

*Skeptical Theism:* The news is full of stories about people losing their homes or their lives in the latest natural disaster. Hundreds of thousands of wildfires burn millions of acres of land each year. But rarely, if ever, do we see how these events fit into the larger picture. We don't know what long-term effect those tragedies had on the victims or those around them. Maybe those events were part of a larger plan. For example, my great grandfather tragically died in an accident when he was only 33 years old. His death had a huge impact on my grandfather (who was 6 at the time) and his future life choices. And it is precisely those life choices that eventually lead him to meet my grandmother. Had he not met my grandmother, I would not exist, and neither would my children or my children's children. I can't say for sure what the long-term impact of that tragedy will be, but what's certain is that the world will be forever different because of that one, seemingly small event. I don't know whether there is a divine being or not, but if such a being does exist, then it has a plan for everything. Those plans, however, are not always going to be visible to us. Our knowledge of how things fit into the bigger picture is so very limited. We also don't know exactly what a divine being would be trying to accomplish. If such a being exists, then it is greater than any of us. It may see reasons for letting something happen that we can't even think of. Perhaps we are like ants trying to

understand the goals and desires of human society. In light of our limitations, it shouldn't be surprising when we don't see any purpose to the evils around us. Maybe their purpose is just hidden from us.

Now, please consider the following story: suppose in a distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering.

DRAFT