

Experimental philosophy and the problem of evil

IAN M. CHURCH, BLAKE MCALLISTER, AND JAMES S. SPIEGEL

Department of Philosophy and Religion, Hillsdale College, Hillsdale, MI 49242

ichurch@hillsdale.edu; bmcallister@hillsdale.edu; jmspieg@gmail.com

Forthcoming in *Religious Studies*.

Please cite published version once it's available; it will be published open access.

Abstract: The problem of evil is an ideal topic for experimental philosophy. Suffering--which is at the heart of most prominent formulations of the problem of evil--is a universal human experience and has been the topic of careful reflection for millennia. However, interpretations of suffering and how it bears on the existence of God are tremendously diverse and nuanced. We might immediately find ourselves wondering why (and how!) something so universal might be understood in so many different ways. Why does suffering push some people toward atheism while pushing others toward deeper faith? What cultural, psychological, or sociological differences account for this diversity of responses? And, importantly, what light might this diversity of responses shed on the problem of evil and how it has been formulated by philosophers in recent years? The primary aim of this paper is to highlight how the tools and resources of experimental philosophy might be fruitfully applied to the problem of evil. In the first section, we review some recent work in this area and describe the current state of this emergent body of literature. In the second section, we review the broader and more recent theoretical developments on the problem of evil. Finally, in the final section, we outline some potential areas of future empirical research that we see as especially promising given those developments.

The problem of evil is an ideal topic for experimental philosophy. Suffering--which is at the heart of most prominent formulations of the problem of evil--is a universal human experience and has been the topic of careful reflection for millennia. However, interpretations of suffering and how it bears on the existence of God are tremendously diverse and nuanced. As empirically minded philosophers, we might immediately find ourselves wondering why (and how!) something so universal might be understood in so many different ways. Why does suffering push some people toward atheism while pushing others toward deeper faith? What cultural, psychological, or sociological differences account for this diversity of responses? And, importantly, what light might this diversity of responses shed on the problem of evil and how it has been formulated by philosophers in recent years?

The primary aim of this paper is to highlight how the tools and resources of experimental philosophy might be fruitfully applied to the problem of evil.¹ To start, in §1, we review some recent work in this area and describe the current state of this emergent body of literature. So far,

¹ Roughly speaking, experimental philosophy is the project of taking the tools and resources of the human sciences (especially psychology and cognitive science) and bringing them to bear on issues within philosophy toward genuine philosophical ends.

experimental philosophy of religion regarding the problem of evil has primarily focused on the two most influential formulations of the problem of evil: William Rowe’s “The Problem of Evil and Some Varieties of Atheism” (1979) and Paul Draper’s “Pleasure and Pain: An Evidential Problem for Theists” (1989). In §2 we review the broader and more recent theoretical developments on the problem of evil. Finally, in §3, we outline some potential areas of future empirical research that we see as especially promising given those developments.

Recent Experimental Philosophical Research on the Problem of Evil

The problem of evil is, of course, an ancient problem. Traditionally, the existence of evil was taken to be logically incompatible with the existence of an all-good, all-knowing, all-powerful God. In his seminal article, “Evil and Omnipotence,” J. L. Mackie put the problem like this: “In its simplest form the problem is this: God is omnipotent; God is wholly good; and yet evil exists. There seems to be some contradiction between these three propositions, so that if any two of them were true the third would be false” (1955, 200). This approach to the problem is called the *logical problem of evil*.

In his landmark book, *God, Freedom, and Evil* (1974), Alvin Plantinga argued that, contrary to Mackie, there is no logical incompatibility between the existence of evil and the claims that God is wholly good and omnipotent. If a world with free moral agents is always more valuable than a world without free moral agents, and if the only way to guarantee a world without evil is to create a world without free moral agents, then a wholly good and omnipotent God might allow for evil to accomplish the goal of having a world with free moral agents.

Plantinga’s argument was broadly taken to be persuasive. In response, proponents of the problem of evil changed tack. Instead of arguing that the existence of evil is logically incompatible with a traditional conception of God, they argue for the much weaker thesis that evil poses an *epistemic challenge* for those forms of theism. Some argued that the existence of gratuitous evil was evidence against traditional brands of theism, a type of argument that came to be called the *evidential problem of evil*. Others framed the problem in terms of epistemic probability—that the amount and distribution of evil in the world made the truth of theism epistemically less probable than alternative hypotheses. This type of argument is called the *probabilistic problem of evil*.² William Rowe’s seminal article, “The Problem of Evil and Some Varieties of Atheism” (1979), is the most cited formulation of the problem of evil in the extant philosophical literature and an archetypal example of an *evidential* problem. And Paul Draper’s landmark article, “Pleasure and Pain: An Evidential Problem for Theists” (1989), has come to be considered, perhaps somewhat confusingly given the title, an archetypal example of the *probabilistic* problem of evil.

As experimental philosophy proliferated at the start of the 21st century, its tools and resources were slow to be adopted by philosophers of religion, and attempts to apply those tools and resources to the problem of evil have only started in the past few years.³ So far, work has exclusively focused on those two seminal formulations, Rowe’s evidentialist formulation and

² It is worth noting that this distinction between *evidential* and *probabilistic* versions of the problem of evil is common but not at all universal. Some scholars simply categorize all versions of the problem of evil that aren’t the logical problem of evil as evidential. We find the distinction between evidential and probabilistic helpful, but there are many ways to carve up the problem of evil pie.

³ Though Helen De Cruz has been a pioneer in the field of experimental philosophy of religion for many years. See De Cruz 2017; 2015b; 2015a; 2018. Other early contributions include Barrett and Church 2013; Leeuwen 2014; Lim 2017; De Cruz and Nichols 2016; and Draper and Nichols 2013. For more on why philosophy of religion was particularly slow to adopt experimental philosophy, see Church *forthcoming*.

Draper’s probabilistic formulation. We will briefly elucidate each before discussing the empirical research that has already been done.

Rowe’s Formulation

In William Rowe’s seminal version of the problem of evil,⁴ he levels the following argument against theism:

1. There exist instances of [*pointless suffering*—defined as] 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)

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.

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

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 the 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).⁵

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, 465–466).

⁴ 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.

⁵ 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.

Empirical Work on Rowe's Formulation

In a new article, McAllister, Church, Rezkalla, and Nguyen propose to study whether the intuitions that underwrite Rowe's seminal formulation of the problem of evil vary across demographics (McAllister et al. 2024, sec. 2.1). They hypothesize that intuitions regarding Rowe's case will significantly diverge according to respondents' religious beliefs, nationality, ethnicity, gender, and education level. McAllister et al. also inquire into the psychological mechanisms underwriting these intuitions. For instance, given that fawns are quintessentially cute animals (think of Bambi!), it raises the question: does cuteness influence the perception of pointlessness? Would the tragic demise of a less adorable animal be considered equally pointless?

Additionally, the concise nature of the fawn case, spanning just two sentences, may play a role in its perceived pointlessness. Some scholars have stressed the significance of context and narrative in addressing the issue of suffering.⁶ Would the inclusion of more background details alter perceptions of pointlessness?

To investigate these questions, Ian Church, Justin Barrett, and the rest of their team conducted an empirical study (see Church, Carlson, and Barrett, 2021; Church, Warchol, and Barrett, 2022), the results of which were more fully reported in McAllister et al.'s 2024 article.⁷ Participants read some variation of Rowe's vignette of the fawn from the 1979 paper. For half of the participants the vignette was accompanied by a description of the role of wildfires in a forest ecosystem, including their contributions to the health of the ecosystem by clearing away dead organic material and leaving behind a topsoil dense in organic materials. The other half of participants read the vignette without context, just as it appeared in Rowe's 1979 paper. The subject of the vignette varied as either a fawn, a boar, or a vulture. Finally, in half of the cases a picture of the subject of the vignette accompanied the vignette. Thus, this experiment contained three variables: context (high or low) picture (picture or no picture) and animal (fawn, boar, or vulture).

After reading the vignette, participants rated the following statements designed to assess their degree of agreement or disagreement with Rowe's intuition that the suffering described in the vignette is pointless: "The story you just read is an example of pointless suffering," "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." Church and the rest of the team initially intended to measure the degree to which participants shared Rowe's intuitions through an index compiled of the score of these three statements, however, they found that whereas scores of the last two questions were highly correlated ($r = .478, p < 0.01$) the first question was not highly correlated in the expected direction with the last two questions ($r = .071, p < 0.01$ and $r = -.171, p < 0.01$). Therefore, they measured agreement with Rowe through an index of the reverse scored second and third questions.⁸

What did they find? In summary, they found significant variations across each demographic variable when it comes to agreement with Rowe, though sometimes in surprising ways—for instance, women were more likely to agree with Rowe than men and agreement with Rowe

⁶ See, for example, Stump 2010.

⁷ The experiment was pre-registered on the Open Science Framework here: <https://osf.io/ebgpd>

⁸ Given that (i) Rowe uses "pointlessness" as a loose shorthand for not bringing about a greater good or preventing a greater evil and (ii) the second and third items correlate with each other but not with the first item, it makes sense to prioritize the second and third items on this index. However, that said, it's worth noting that using the three-item index or the just first item would not radically change McAllister et al.'s results.

consistently dropped with increases in education-level. They also found that the inclusion of context dramatically diminished the perceptions of pointlessness. The cuteness of the animal (with or without a picture) did not have a measurable impact on people's responses. Perhaps most remarkable was the fact that so few people actually agreed with Rowe's assessment of the target vignette. To quote McAllister et al., "One of the more notable results is how few people, only 12.61%, share any level of agreement with Rowe about the gratuitousness of the fawn's suffering." (2024, sec. 3.1).

McAllister et al. argued that, if these results are accurate, then they put those who agree with Rowe in an epistemically perilous position. Few, it seems, will have strong reasons for privileging the judgment of those who agree with Rowe over those who don't. Accordingly, one's justification for believing that the fawn's suffering is gratuitous could be partially or even fully defeated, undermining Rowe's argument. Given the exploratory nature of the research thus far, McAllister et al. acknowledge that such a conclusion is highly defeasible and conditional on whether the study is accurately measuring agreement with Rowe. That said, these findings certainly raise questions about the reach and strength of Rowe's argument—questions that deserve further investigation.

Follow-Up Research on Rowe's Formulation

Serious questions remain about the accuracy of these findings and what to make of them. To start, there is the question of whether McAllister et al.'s 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 by the fawn's suffering (or some worse evil prevented by it) but think of this possibility as extremely remote. To disagree with Rowe, however, participants need to think of this possibility as being plausible, or at least not particularly unlikely.

To address this concern, McAllister, Church, and the rest of their team have developed an eight-item index to measure agreement with Rowe which they plan to use in follow-up research. First, the two from McAllister et al. (2024):

"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 six 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*)

This expanded index should allow them 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.

In their 2022 article, “The Context of Suffering,” Church, Warchol, and Barrett draw on this empirical research to explore the idea that context plays an important role in diminishing perceptions of suffering as pointless. Referees and critical commentators raised a suite of important questions about *why* and *whether* context really diminish perceptions of pointlessness and what significance this “contextual effect” has for evaluating Rowe’s argument.

One suggestion offered by commentators 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, as some referees mused, the context that was given to participants contained a hidden theodicy. On this proposal, contexts drive down the perceptions of pointlessness only (or primarily) when they invoke a theodicy. Call this the *sneaky theodicy hypothesis*. Or we might think that the context is simply distracting participants from the pointlessness of the suffering. So on this hypothesis, it’s not as if participants are somehow seeing a “point” in the suffering; instead, the sense of 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 it is this awareness that is driving down perceptions of pointlessness. Call this the *skeptical theism objection*.

In follow-up research, McAllister, Church, and the rest of their team aim to explore each of these hypotheses by introducing new context conditions. Let’s call the original two context types “Regular Context” and “No Context.” The following three context types will be added in future research: “Negative context,” “Theodicy Context,” “Junk Context” to address the first three of those objections.

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

Addressing the sneaky theodicy hypothesis, the “Theodicy Context” will embed a full-throated theodicy into the original vignette. If the sneaky theodicy hypothesis were true, we would expect to see the “Theodicy Context” significantly diminish the perception of pointlessness, even more than the “Regular Context” condition. If the “Theodicy Context” doesn’t diminish perceptions of pointlessness as significantly as “Regular Context,” then we may doubt that the sneaky theodicy hypothesis best explains the observed contextual effect.

Addressing the distraction hypothesis, “Junk Context” will provide 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 context significantly reduces the perception of pointlessness, then that would lend further credence to the distraction hypothesis. If “Junk Context” doesn’t significantly diminish participants’ perception of pointlessness, then it seems unlikely that the distraction hypothesis best explains the contextual effect.

Finally, in response to the skeptical theism hypothesis, instead of being assigned to one of the five context conditions, some participants will be randomly assigned to what we’re calling the Skeptical Theism condition. Participants who are assigned this condition will 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*, *sneaky theodicy*, *distraction*, or *skeptical theism*--best explain the role context played in diminishing perceptions of pointlessness. Further research could suggest that the role context is playing is more mysterious than we might have previously expected.

Draper's Formulation

In formulating his probabilistic problem of evil, Draper does not argue that God's existence is logically incompatible with evil but only that "our knowledge about pain and pleasure creates a problem for theists" (1989, 331). Draper asks us to let "O" represent "a statement reporting both the observations one has made of humans and animals experiencing pain or pleasure and the testimony one has encountered concerning [such things]" (1989, 332). He then contends that O is less epistemically probable given theism than the *hypothesis of indifference* (HI).⁹ He provides the following definitions:

HI: "Neither the nature nor the condition of sentient beings on earth is the result of benevolent or malevolent actions performed by non-human persons." (1989, 332)

Theism: "There exists an omnipotent, omniscient, and morally perfect person who created the Universe." (1989, 331)

Let "P(X|Y)" stand for the epistemic probability that X given Y. We can now express Draper's conclusion a bit more formally as:

Draper's Conclusion: "P(O | HI) is much greater than P(O | theism)" (1989, 333)

Note that according to Draper, the epistemic probability of O given HI isn't just *greater* than the epistemic probability of O given theism; it's *much greater*. Hence, the evidence that O provides for HI over theism is supposedly quite significant.

Importantly, the exact degree to which O confirms HI over theism will depend on each subject's unique epistemic situation. For one, not all subjects start with the same "O" since not everyone has the same observations or testimony about pleasure and pain. For another, background information may differ from one person to another such that the relevant conditional probabilities will also differ from subject to subject.¹⁰ Nevertheless, Draper seems confident that his conclusion will hold for most if not all subjects given "the biological role played by both pain and pleasure in goal-directed organic systems" (334). Such a role fits perfectly with the hypothesis of indifference but not nearly so well, he claims, with theism.

⁹ Draper provides the following analysis of epistemic probability:

Relative to K, p is epistemically more probable than q, where K is an epistemic situation and p and q are propositions, just in case any fully rational person in K would have a higher degree of belief in p than in q. (1989, 349n)

¹⁰ Background information, if not mentioned explicitly, is always tacitly featured in conditional probabilities. That is, P(O|theism) should really be written P(O|theism & B) where B represents the subject's background evidence. The present point is that two subjects can have two different sets of background evidence such that P(O|theism & B₁) may differ from P(O|theism & B₂).

In these early stages, experimental philosophers have primarily been interested in exploring whether people's own assessments of the relevant conditional probabilities align with those of Draper. Ian Church, Blake McAllister, James Spiegel, and Justin Barrett (hereafter Church et al.) explore this in their manuscript, "Testing Folk Perceptions of Draper's Probabilistic Problem of Evil." Their initial hypothesis was that agreement with Draper's Conclusion—i.e. answers indicating that one's "O" is more epistemically probable given HI than given theism—would *not* be widely shared across demographic variables.¹¹ If anything, in keeping with the empirical research on Rowe's formulation discussed above, Church et al. hypothesized agreement with Draper's Conclusion to be in the clear minority.

They also expected atheists and agnostics to be significantly more likely than various religious groups (e.g. Christians, Hindus, etc.) to give answers in line with Draper's Conclusion. Likewise, given that men are statistically more likely than women to be atheists (Cragun 2016, 307), they expected men to agree with Draper's Conclusion significantly more on average than women. And given that education level negatively correlates with religiosity (Beit-Hallahmi 2006, 313), they expected agreement with Draper's Conclusion to correlate with education level.

To gauge agreement with Draper's Conclusion, participants were given the following two thought experiments in a random order:

Planet X: Imagine that you are transported to another universe, to a planet (call it Planet X) that is very similar to Earth and with creatures (including humans) that are very similar to Earth's creatures and humans. Before exploring Planet X, imagine that you also know that this planet was created by an all-powerful, wholly good God who is involved in the lives of the humans there. Based solely on this knowledge, how confident are you that Planet X would contain a similar distribution of pain and pleasure that you have observed (and heard about from other people) in our own world?

Planet Y: Imagine that you are transported to another universe, to a planet (call it Planet Y) that is very similar to Earth and with creatures (including humans) that are very similar to Earth's creatures and humans. Before exploring Planet Y, imagine that you also know the condition of the humans on Planet Y is NOT the result of actions performed by supernatural beings. Based solely on this knowledge, how confident are you that Planet Y would contain a similar distribution of pain and pleasure that you have observed (and heard about from other people) in our own world?¹²

Using a 7-point Likert confidence scale ranging from "not at all confident" to "entirely sure," participants were asked to rate their confidence that the given planet would have a "similar distribution of pain and pleasure as [they] have observed (and heard about from other people) in our own world." Plausibly, the more confident one is that the pleasure and pain on Planet X or Y will resemble our own, the higher one takes $P(O | \text{theism})$ and $P(O | \text{HI})$ to be, respectively.

¹¹ Work in experimental philosophy has repeatedly shown that many (though certainly not all) philosophical intuitions can vary based on education, gender, ethnicity, etc., and we expected to find a similar result here. See, for example, Weinberg, Nichols, and Stich 2001; Machery et al. 2017; Stich and Machery 2022. For work on the stability of philosophical intuitions across cultures, see Knobe 2019.

¹² This research was pre-registered on the Open Science Framework. See here: <https://osf.io/xqtn5>

Given this, Church et al. defined agreement with Draper's Conclusion as registering a Planet Y confidence score (ranging from 1-7) that is greater than a Planet X confidence score (again, ranging from 1-7). Falling under this umbrella are two different types of agreement:

Partial Agreement: A participant reports partial agreement with Draper's Conclusion if and only if their Planet X score is lower than their Planet Y score by one point.

Full Agreement: A participant reports full agreement with Draper's Conclusion if and only if their Planet X score is at least two points lower than their Planet Y score. (Church et al., n.d.)

Contrary to what Church et al. hypothesized, their research found that Draper's Conclusion enjoys comparatively broad intuitive support amongst the folk. Where only 12.6% of participants "share any level of agreement" with William Rowe's 1979 formulation of the problem of evil (McAllister et al., 2024, sec. 3.1). Around 41% of participants surveyed by Church et al. reported at least Partial Agreement with Draper's Conclusion. 28% registered Full Agreement. Put another way, this line of research suggests that Draper faces far less of an uphill battle regarding folk intuitions than Rowe does. Draper took his conclusion to capture the intuitions "of a great many people who have regarded evil as an epistemic problem for theists," and he certainly seems to be correct in that regard (1989, 332).

Church et al. did find some variation across demographics, but variations across gender, ethnicity, nationality, education level, income, etc. were far more muted than the variation found in the McAllister et al. study of Rowe's argument. Religion, however, seemed to have an enormous influence on whether or not (and to what degree) participants registered agreement with Draper. Atheists and agnostics were significantly more likely ($p < .001$) to agree with Draper than any of the other religious categories considered (such as Protestant, Catholic, Other Christian, Hindu).

Across all participants, the mean confidence score for Planet Y was higher ($M = 4.83$, $SD = 1.44$) than Planet X ($M = 4.41$, $SD = 1.72$). A 1-tailed t-test showed that the differences between the Planet Y scores and Planet X scores was significant ($p < .001$; $d_s = .26$). Other things being equal, Church et al. propose that there may be collective wisdom in the aggregate scores—that other things being equal, we should expect the average confidence scores of all participants to be a more accurate measure of the conditional probabilities relevant to Draper's Conclusion than the individual confidence scores of any (arbitrarily chosen) individual participant. If that's right, then such findings might provide evidential support for Draper's Conclusion.¹³

Follow-Up Research on Draper's Formulation

Church et al. caution that further research is needed before any firm conclusions can be drawn. To start, Church et al. acknowledge that atheists and agnostics are significantly over-represented in their sample. Given that atheists and agnostics were significantly more likely to agree with Draper than other religious groups, this over-representation may have skewed the averages in Draper's direction. Church, McAllister, and the rest of their team are conducting follow-up studies with a more representative sample.

Another worry arises from the way Church et al.'s study directly asks participants how confident they are that "an all-powerful, wholly good God" would create a world with a similar

¹³ And that would be evidential support in addition to the support Draper provides from his "armchair" argument.

distribution of pain and pleasure as our own. Many theists, we might worry, will be quite confident that God would create such a world, because they already believe that God created *this* world. This is, however, to stack the deck in favor of theism and so might be called the *improper weighting effect*. Their evaluation of $P(O | \text{theism})$ *should* be conducted given only the fact that an all-powerful, wholly good God exists (and general background evidence acceptable to all parties).

A related concern is that the explicit religious framing might tempt participants to simply parrot whatever they take to be theological (or atheological) orthodoxy. Call this the *parroting effect*. Given the significant variation along religious lines, we might worry that reported answers aren't really reflecting the intuitions of participants but merely their pre-established theological or atheological commitments.

To address both of these concerns, Church's research group is designing follow-up studies that discourage improper weighting or parroting by reorienting the questions around third-parties (e.g., "What should this *other person* with limited information think?") or removing explicitly religious framing (e.g., "What would this world look like if governed by a wholly-good AI?").

A final worry is that Draper's exclusive focus on pleasure and pain may limit the reach of these results. People may agree that the distribution of pleasure and pain fits better with the hypothesis of indifference than with theism, but would they say the same about good and bad states of affairs more generally? What about distributions of virtue and vice? Beauty and ugliness? Once again, Church's team plans to explore these issues further.

Philosophical Research on the Problem of Evil: The Current State of Play

The literature on the problem of evil is immense, and extends far beyond the work of Rowe and Draper. There is no hope of us surveying it in its entirety. Our overview will be selective, giving shape to the broad contours of the debate but with a preference for certain issues that we find especially suitable for empirical probing, which will be discussed below..

Responses to Evidential Formulations

Non-empirical responses to the evidential and probabilistic problems of evil have generally taken three forms. The first is to give a *theodicy* (or a defense). The second is to adopt *skeptical theism*. The third, only now gaining prominence on the contemporary scene, is to claim for God some form of *moral immunity*.

Both theodicies and defenses articulate reasons that would justify God in allowing evil (or certain types of evil) to exist. A defense claims only that it is epistemically possible, or true for all we know, that God has such reasons. A theodicy claims it is plausible God has such reasons.¹⁴ A theodicy is more suitable in this context, since the mere possibility of justifying reasons does not seem sufficient to defuse either Rowe's or Draper's arguments.¹⁵ Skeptical theism contests our ability

¹⁴ Plantinga 1974 defines theodicies as maintaining that God *does* have such reasons. We are instead following the usage of van Inwagen 2006 and others.

¹⁵ One can acknowledge the possibility of justifying reasons and still reasonably believe that there are none. Hence, a defense may not block reasonable belief in premise 2 of Rowe's argument. As for Draper's, the mere possibility of justifying reasons may raise the probability of O given T slightly but not enough to prevent significant disconfirmation of theism.

to tell whether there are any justifying reasons.¹⁶ In particular, we are in no position to conclude that God does not have justifying reasons. A final strategy says there is no need for justifying reasons since God enjoys moral immunity. That is, whatever moral principles underlie the contention that a good God would not allow evil (or evil of a certain sort) do not apply to God. We'll look at each type of response more closely, from last to first.

Moral Immunity

The terminology of “moral immunity” is our own and seeks to cover a variety of positions all united by the contention that none of the moral principles on which the problem of evil is grounded, nor any other moral principles in their general vicinity, apply to God. Some go so far as to claim that God is not morally good (Davies 2006). He does not exist in our moral economy in such a way that requires justifying reasons for acting in certain ways. Indeed, God does not act for reasons at all. Others concede God's moral goodness but deny that his goodness gives him any obligation to promote our well-being (Adams 1999; Murphy 2017). For Mark Murphy, God has reasons to promote our well-being but not *requiring* reasons. It may still be the case on such views that God's treatment of us must meet certain requirements if he is to count as good or loving toward us (Adams 2013, 16-17), but whether he is good or loving toward us is up to his discretion. Still others might concede that God has moral obligations to us but contest John Stuart Mill's insistence that these must be substantially on par with our own (Harrison 2017). Treatment that may be forbidden for us can be permissible for God. While growing in popularity, such criticisms of the problem of evil remain in the minority, perhaps in part due to the predominance of Christianity within the debate.

Skeptical Theism

Much more common is skeptical theism. Skeptical theists attack key inferences or rational processes relied on to support the problem of evil. For instance, how does one support the notion of gratuitous evils? Presumably, by surveying all of the justifying reasons we can think of and, finding none of them satisfactory, inferring that there are no justifying reasons. Skeptical theists argue that, in this context, such *noseum inferences* should not be trusted.

Skeptical theists substantiate their skepticism in a variety of ways, all of them highlighting reasons why we should take ourselves *not* to be in a position to discern potential justifying reasons. In one of the earliest defenses of skeptical theism, Stephen Wykstra (1984) appeals to a parent-child analogy. Just as we should not expect a small child to understand why her parents allow her to suffer, even when there are such reasons (e.g., she is receiving cancer treatments), so we should not expect ourselves to understand why God allows us to suffer. Appeal to the parent-child analogy is controversial and raises larger questions about what responsibilities God has to (try to) communicate his plans to us and/or to comfort us with his presence in the midst of indecipherable suffering. Here discussions of the problem of evil and the problem of divine hiddenness converge.

Alston (1991) defends skeptical theism more systematically by inventorying limitations of the human cognitive condition vis-à-vis potential justifying reasons:

1. Lack of relevant data.
2. Complexity greater than we can handle.

¹⁶ Skeptical theism might better be titled “agnosticism,” as Howard-Synder 2009 suggests, since it can be endorsed by theists and non-theists alike.

3. Difficulty of determining what is metaphysically possible or necessary.
4. Ignorance of the full range of possibilities.
5. Ignorance of the full range of values.
6. Limits to our capacity to make well-considered value judgments. (Dougherty 2014)

If we are limited in these ways, then it is difficult to see how we could ever reasonably believe there to be no justifying reasons for allowing evil.¹⁷ The most pressing challenge for skeptical theism is usually to prevent their skepticism from spilling over into unwarranted skepticism in other domains (moral, religious, perceptual, etc.).

Traditional Theodicies

Theodicies try to articulate plausible justifying reasons for God to allow evil. They tend to fall into a few basic categories. Free will theodicies, which trace back at least to Augustine,¹⁸ start with incompatibilism—the view that a choice cannot be both free and determined. Thus, not even God can allow creatures to make free choices and ensure that they never use that freedom to produce moral evil (evils caused by free will). Nevertheless, God is justified in allowing creatures to make free choices, and whatever moral evil may come of them, because doing so is necessary to make possible great goods such as moral righteousness or genuine loving relationships. Contemporary defenders often stress the importance that our free choices have more than token significance. Plantinga (1974) emphasizes the need for “morally significant freedom” while Swinburne (1998) touts the value of bearing responsibility for our own welfare and that of others. Greater good theodicies, in contrast, argue that evil is integral to achieving a greater good. John Hick’s (1976) soul-making theodicy is the most prominent of this sort. Hick argues that the only environment suitable for growth in virtue is one in which suffering is present. Without it, people would lack the motivation and opportunity to acquire and exercise moral characteristics such as courage, compassion, generosity, selflessness, and the like.

Free will and greater good theodicies are often presented as rivals, and there are certainly important differences between them. In free will theodicies, evil is not necessary. If humans always used their free will rightly, then we could achieve the desired goods without any evil. Evil is an avoidable byproduct. Whereas in greater good theodicies, evil is necessary for (or at least integral to) producing the great goods in question. That being said, almost all free will theodicies *also* make appeal to greater good considerations, and many greater good theodicies *also* rely on incompatibilist conceptions of free will.¹⁹ Regarding the former, those advocating for free will theodicies almost always add that God will not permit the creation of free creatures to be an unmitigated tragedy.²⁰ Thus, God ensures that the moral evil resulting from free choices is ultimately used for good. Regarding the latter, many of those advocating for greater good theodicies rely crucially on the idea that free choices cannot be determined. For example, Hick’s account of why suffering is essential to soul-making is only plausible on the supposition that God cannot develop these virtues within us (or cannot do so in an equally valuable way) without our free and undetermined cooperation.²¹

¹⁷ Most other defenses of skeptical theism build a case for one or more of the limitations identified by Alston. Notable examples include van Inwagen 2006 and Bergmann 2001.

¹⁸ See Augustine’s *On Free Choice of the Will*.

¹⁹ For more on this point, see Spiegel 2011.

²⁰ Davis 2001 provides a good example of this.

²¹ See Hick 1976, 255-256. See also Eckstrom 2021, 33-34.

So it may be that there is less tension between these two approaches than commonly supposed. Be that as it may, not all greater good theodicies assume incompatibilism. Compatibilist-friendly theodicies include the divine glory theodicy and the “O Felix Culpa” theodicy. The former says that the existence of moral evil allows God to display his mercy and love in saving sinners and his justice and wrath in punishing them, thereby increasing his own glory.²² The latter says that the very best worlds are those in which God is allowed to execute his rescue plan of becoming incarnate, atoning for sin, and raising fallen creatures to new life. However, such a plan can be executed only in a world filled with moral evil.²³

So far, most of these theodicies have focused exclusively on moral evils—those stemming from the misuse of free will. But what of natural evils? These are evils that do *not* result from free will such as earthquakes, diseases, or non-culpable accidents (e.g., skidding on a patch of ice while driving resulting in the death of the driver). Some theodicies, such as Hick’s soul-making theodicy, are equipped to explain natural evils as well as moral evils, but often multiple theodicies must be patched together to account for the full variety of evils in the world. Regarding natural evils, a prominent theodicy is the natural law theodicy. This maintains that any world with sufficiently stable natural laws will inevitably contain natural evils, and that the goods made possible by stable natural laws (such as providing a suitable environment for the exercise of human freedom and responsibility, the development of human moral character, and the pursuit of scientific discovery) are enough to justify the allowance of such evils.²⁴

Traditional Challenges for Theodicies

Theodicies face several perennial challenges. First, we have already seen the challenge in accounting for the full variety of evils observed in our world. Even if successful, a theodicy may justify one type of evil (e.g., human suffering) while leaving another untouched (e.g., animal suffering). Second, many of the most plausible theodicies rely crucially on incompatibilism about free will, or at least, on the idea that libertarian freedom is significantly more valuable than whatever sorts of freedom are compatible with determinism. This assumption is highly controversial, however.²⁵ Third, the plausibility of a theodicy often comes down to the question, “Is it worth it?” Are the goods to which the theodicy appeals *good enough* to justify the evils in question? This is a basic value judgment.²⁶ What is the theodicy to do with those (seemingly always a significant number) who reasonably disagree? Fourth and finally, many intuit that allowing evils of an especially horrendous sort is inherently wrong. In other words, allowing those evils is itself an evil and, as the Pauline Principle tells us, one must never do evil that good may come of it.²⁷

The Rise of “Good-Hearted” Theodicies

²² This theodicy is most prominent among Calvinists. See, e.g., some of the essays in Johnson and Alexander 2016 (especially those of Johnson, Green, and Hart).

²³ Plantinga 2004. While this is a compatibilist friendly theodicy, it doesn’t *require* compatibilism either. Plantinga himself is an incompatibilist.

²⁴ Reichenbach 1976, Swinburne 1998, van Inwagen 2006, and Hasker 2008, Ch. 5.

²⁵ See, e.g., Eckstrom 2021 and Speaks 2015, 97-98, on this point. According to the PhilPapers 2020 survey, 57.68% of respondents accepted or leaned toward compatibilism compared to only 18.20% who accepted or leaned toward libertarianism.

²⁶ See, e.g., Hick 1976, 255-256.

²⁷ Sterba 2019 appeals to this Pauline Principle (taken from Romans 3:8) in supporting his version of the problem of evil.

Work on theodicy reached something of a crescendo near the turn of the millennium, with a flurry of seminal books and articles being published by William Hasker, Alvin Plantinga, Richard Swinburne, Peter van Inwagen, and others. Among the most innovative was Marilyn McCord Adams' *Horrendous Evils and the Goodness of God* (1999). This book, along with earlier work by Eleonore Stump (1985), seems largely responsible for the subsequent rise in what Evan Fales calls "good-hearted theodicies" (Fales 2013, 352), the most notable of which is found in Stump's *Wandering in Darkness* (2010).²⁸

Following the lead of Adams and Stump, good-hearted theodicies emphasize God's love for all individuals and insist that this love be expressed in ways we can appreciate when viewed in their proper context. That is, good-hearted theodicies try to ensure that the charitable reader will not revile God's decisions nor see him as indifferent toward human suffering. This concern for maintaining a sense of God's compassion and care for individual persons is perhaps why Fales also calls such theodicies "soft-hearted" (Fales 2013, 351-352). This sensitivity is manifest in several commonly proposed requirements on the content of good-hearted theodicies, each of which is contentious among theists more generally:

1. *The Horrendous Evil Requirement.* A theodicy must account not only for evil in general but for the most horrifying instances of it.
2. *The Good Life Requirement.* A life that is worth living (on-the-whole good and meaningful) must be guaranteed, or at least made available, to all persons, including those who endure horrendous evils.²⁹
3. *The Sufferer-centered Requirement.* Horrendous suffering must result in some sufficiently great good for the one who suffered, lest he or she be treated as a mere means to an end.
4. *The Defeat Requirement.* The sufficiently great good that accrues to the sufferer must arise organically from the suffering itself, thereby "defeating" the suffering (Chisholm 1968) and avoiding any impression of an arbitrary pay-off.
5. *The Redemption Requirement.* The defeating good must be subjectively appropriated by the sufferer such that he or she no longer wishes away that suffering.

Along with these substantive constraints come certain stylistic ones as well. Proponents try to write with empathy and are careful not to downplay the horror of suffering.

In seeking to meet all of these requirements, several important trends have emerged, each representing a key development in the domain of theodicy.

1. *Eschatology.* It is plain that defeat and redemption are not always made available in this life, so good-hearted theodicies have followed Hick (1976, 338-341) in relying crucially on eschatological goods.

²⁸ Adams herself eschews the word "theodicy" and does not attempt to give a justification of God in the traditional sense. She does, however, attempt to reconcile the goodness of God with evils in our world (Adams 1999), and many of the unique aspects to her approach can and have been incorporated into theodicies by Stump and others.

²⁹ The debate over whether the goods under consideration must be guaranteed (as in Adams 1999) or only made available (as in Stump 2010) applies to 3-5 as well, though I will ignore this nuance for the sake of readability. Some only require that this good be *made available* to the sufferer (and likewise for the goods in 3-4).

2. *Special Revelation*. In fleshing out an eschatological vision, good-hearted theodicies have found it helpful, if not indispensable, to make use of the resources of particular religious traditions, most commonly Christianity.
3. *Worldviews and Perspectives*. As the appeal to specific religious traditions makes clear, in giving a theodicy one can no longer assume that the larger worldview supposed by that theodicy is shared by (or even comprehensible to) most of one's readers. In this case, the theodicy is liable to be misunderstood and unfairly dismissed (Stump 2010). It therefore becomes incumbent on the theodocist to articulate this larger worldview and make it as intuitive as possible. Hence, how things seem to the individual—i.e., one's "perspective"—becomes an increasing point of focus (McAllister 2020).
4. *Narrative*. Some argue that embedding a theodicy within a narrative gives the reader special understanding of that theodicy and the worldview presumed by it—one that registers on an intuitive and emotional level (Stump 2010, Ch 2-4 and 373-374).
5. *Pastoral Awareness*. While not completely disowning Plantinga's distinction between the philosophical and the pastoral (Plantinga 1974, 63-64), good-hearted theodicies seem aware that pastoral considerations cannot be altogether ignored in one's philosophical response to the problem of evil (McAllister 2020, 445).

Based on current trends, we think the future of theodicy most likely lies in this good-hearted direction.

Potential Areas of Experimental Philosophical Research on the Problem of Evil

Having surveyed the most prominent formulations of the problem of evil as well as the range of responses to the problem, let's now consider how experimental philosophy might continue to contribute to that emerging scholarly discussion. Below is just a sampling of some of the lines of research that come to mind. There are certainly many more.

As regards the moral immunity approach to the problem of evil, it would be worth empirically exploring how people think about the goodness of God. Do people think of God as being morally good? If so, then do they think that God's moral goodness and obligations are similar to those of humans? And what correlations might there be between various theists' answers to these questions and the doctrinal commitments distinctive to their respective theological traditions? How is someone's understanding of divine goodness impacted by the images of God they're most attracted to (see, for example, Johnson et al. 2019)? Could it be that those who rely more heavily on religious coping strategies think of God's goodness as more similar to our own, perhaps as having moral attitudes and obligations toward us similar to that of human parents?³⁰ Could those who think more analytically, or who score lower in cognitive empathy, be more inclined to think of God's goodness as non-moral or entailing little by way of beneficence? Answers to such questions might be helpful not only in ascertaining why particular theists are or are not drawn to the moral immunity approach, but may even bear on how rational we find those approaches and/or rejections of them.

The skeptical theist approach likewise raises many questions warranting experimental philosophical inquiry. For starters, why do people believe that certain forms of evil are gratuitous? Are such convictions bare intuitions or are they conclusions reached on the basis of inferences? In

³⁰ Freud is wrong that we believe in God's existence because we want a daddy in the sky, but maybe some of us believe God is *morally good* for this reason. Work applying attachment styles to perceived relationships to the divine might be especially relevant here. See, for example, Bradshaw, Ellison, and Marcum 2010.

either case, what sorts of cognitive processes are involved? Skeptical theists tend to appeal to the “noseeum” inference, but is this the true basis of most people’s beliefs? If it’s not, then skeptical theists may be attacking the wrong thing.

Regarding the plausibility of skeptical theism itself, how credible do the folk find the notion that we really are in the dark as regards the potential justifying reasons for God’s allowance of evil? What grounds might they have for this skeptical stance?³¹ Is this attitude prompted by rational reflection, personal experiences, the influence of their religious tradition, or something else? Do they display the same level of skepticism in other, non-religious contexts? If not, is this because they are being more careful here or because they are engaged in motivated reasoning to preserve prior religious commitments? We might also wonder if many people have inconsistent beliefs regarding their ability to discern gratuitous evils. Perhaps, for example, many people believe that their inability to see a “point” to an instance of suffering counts as strong evidence for thinking there is no point, while also confidently affirming that they wouldn’t be privy to a “point” if there was one (see Church, Carlson, and Barrett 2020). The answers to such questions will help us better understand the draw to skeptical theism and might even increase or diminish how much credence we assign it.

Regarding theodicies, a wide range of issues call for experimental probing. First off, where do the folk generally land as regard preferences for either the free will or greater good theodicy? And how might these preferences be correlated with people’s doctrinal commitments about such things as divine providence and the authority of Scripture? How might people’s theodicy preferences be correlated with their intuitions about God’s priorities for humanity? What are people’s basic value judgments regarding what evils God could reasonably allow and for what reasons (e.g., the goods of respecting our moral autonomy, increasing our moral maturity, making genuine relationships possible, the achievement of great beauty, etc.)? In the view of the folk, is the display of particular divine attributes reasonable grounds for God’s permission of evil? If so, what attributes do people find most compelling in this regard (e.g., divine justice, divine mercy, divine loving relationality)?³² How do people’s assessments differ, if at all, regarding their assessments of God’s permission of natural and moral evils? And what sorts of factors do people regard as potential grounds for divine permission of horrendous evils? Is certainty of full and final redemption necessary, sufficient, or both? What are the folk intuitions regarding whether a loving human parent would bring a child into the world knowing that he or she would eventually experience horrendous evil? And do these intuitions line up with their intuitions on whether God would do the same? What sorts of considerations might inform those for whom these intuitions do not correlate?

With regard to such questions, the point is not simply to understand *what* the folk think but *why* they think it. What are the underlying cognitive, psychological, social factors influencing their judgments and what insight might this give us into the nature of the debate and (possibly) into the plausibility of various positions within that debate?

To give an example, moral foundations theory has suggested that framing the moral landscape exclusively in terms of harm/care and fairness/reciprocity is a hallmark of secular liberalism, whereas religious conservatives tend to draw from a more diverse moral palate (Haidt 2009). Since many theodicies appeal to considerations beyond harm and fairness, could this division help explain why people assess the plausibility and prospects of theodicies so differently? This would certainly make sense of why the problem of evil ascended to new heights alongside the rise of

³¹And we might wonder if an overly confident assessment of an evil’s gratuitousness might be a product of the Dunning-Kruger effect.

³²And we might plausibly wonder if the kinds of theodicies people find most intellectually attractive are also the ones they find pastorally or even existentially satisfying.

secular liberalism in the modern West,³³ as well as why those problems are nowadays almost exclusively formulated in terms of suffering and well-being (or, as with Draper, pain and pleasure). It might also explain why good-hearted theodicies, which give special focus to considerations of harm/care and fairness/reciprocity, are gaining in prominence.

Speaking of this, the recent movement toward good-hearted theodicies also suggests a gamut of questions for experimental inquiry. What are we to make of the increasing focus on horrendous suffering? Is this a constructive development (honing in on those cases that pose the greatest rational challenge to God's existence) or a distracting rhetorical tactic (rife for triggering emotional reactions and one-sided judgments)?³⁴ Could it perhaps be both? To what extent do people find sufferer-centered requirements plausible, and why or why not? As floated above, does their appeal stem from a devoted concern to individual well-being and fairness most common to the modern West? In any case, do theodicies that meet such requirements stand a better chance at persuading audiences today than those that don't? When it comes to the final defeat or redemption of evil, what sorts of factors do people find most compelling (e.g., moral maturity and Godlikeness, the experience of divine gratitude or love, etc.)? Is appeal to an afterlife necessary for final redemption to be plausible? If so, must one inevitably approach these issues from the standpoint of a particular theological perspective beyond that of bare theism? How do one religion's assets for dealing with the problem of evil compare to those of alternative religions? And how does the strength or weakness of various faith traditions in this regard speak to their ultimate sustainability in light of the challenges posed by evil?

Still other important questions related to the problem of evil more generally await probing, particularly pertaining to the various presuppositions people bring to their reflections on God and evil. As should already be plain from our discussion thus far, an especially relevant area to probe is moral psychology. While it can easily be overlooked, assessments of the problem of evil are ultimately moral in nature, involving basic judgments about what an all-powerful, perfectly good God *ought* or *ought not* to do, usually based on an evaluation of whether certain goods are *good enough* to outweigh or defeat the evils that make them possible. Obviously, then, we might expect one's moral-theoretical commitments and tendencies—even if only inchoate or intuitively expressed—to affect one's take on the problem of evil. But how exactly does this play out? For example, those on both sides of the problem of evil have accused the other of overly-utilitarian thinking.³⁵ Do we see those with more consequentialist intuitions tending toward one side of the argument or the other? The same might be asked of those whose intuitions are more deontological or virtue-oriented.

One may also easily overlook the fact that such moral assessments are necessarily influenced by one's own moral-psychological condition. So, for starters, we might wonder how one's level of moral seriousness impacts one's take on the problem of evil? Could there be a correlation between how committed a person is to living rightly and how inclined one is to make particular kinds of judgments about whether God would allow evil, or the sorts of evils, we find in this world? To what degree and in what way do people's hedonistic tendencies impact their assessment? In terms of

³³ Ancients and medievals tended to think of the problem of evil differently—often as more of a metaphysical or ontological problem than a moral one—and generally found it much less convincing.

³⁴ Returning to moral foundations theory, Haidt 2012, 154-155, proposes that our harm/care moral foundation originated as a way of encouraging concern for vulnerable children. Is it any wonder, then, that the examples of horrendous evils that add the most credence to the problem of evil (such as those presented by Ivan in Dostoyevsky's *The Brothers Karamazov*) are ones featuring children?

³⁵ Critics of the problem of evil, particularly greater good theodicians, are regularly accused of employing “ends justify the means” reasoning. Proponents of the problem of evil are sometimes accused of thinking that God is interested only in creating a hedonistic paradise.

personal psychology, we might ask how people's level of emotional empathy affects their evaluation of the argument. Or, looking at personal psychological characteristics more broadly, how might whole personality types differ in their approaches to the problem of evil? Empirical studies could be conducted using any of the major personality typologies, such as the OCEAN model or "Big Five" personality inventory (Goldberg, 1993) and the HEXACO model of personality structure (Lee and Ashton, 2004).

In justifying God's allowance of evil, theodicians most often appeal to moral or spiritual goods, but some have leaned heavily on aesthetic goods as well.³⁶ Since people vary not only in taste but in sensitivity to beauty and interest in aesthetic considerations more generally, a study probing the relationship between such factors and judgments about the problem of evil could be telling. Could the plausibility of certain theodicies vary alongside aesthetic sensibilities or tastes? Perhaps those drawn to epic storytelling are more positively inclined toward "O Felix Culpa" theodicies, which appeal to the beauty and glory of forgiveness and redemption. Perhaps only those without a strong concern for aesthetics find pure free will theodicies (which cast moral evils as unmitigated tragedies) persuasive. We would not be surprised if many theodicies are rejected due to a failure to adequately portray the ugliness of evil as being organically integrated into an aesthetically pleasing whole. This would explain why the defeat of evil has become an increasing prominent element to theodicies, since the notion of defeat is drawn originally from aesthetics.³⁷ In any case, we suspect that the beauty or ugliness of a theodicy and the story it tells plays a larger factor in people's acceptance or rejection of that theodicy than is commonly appreciated.

The problem of evil hinges not only on moral and aesthetic assessments but on epistemological ones as well. Thus, we might find it fruitful to explore what personal psychological factors underlie these assessments. One basic question concerns how much epistemic self-confidence people have, both in themselves and in the effective reach of the human intellect more generally? Do they tend to trust their intuitive judgments or are they more prone to reflective scrutiny and second-guessing? How comfortable are people with admitting ignorance and acknowledging mystery on these issues? Do they display a strong need for closure or struggle with uncertainty? How ready are they to defer to "the experts" when it comes to the problem of evil? Who even do they perceive as experts on these issues (philosophers or religious leaders)? And how does all of this connect with their willingness or unwillingness to rely on faith?

With respect to both epistemic and axiological evaluations, we might also wonder how broader cultural factors are biasing, steering, or otherwise influencing our thinking about the issue, whether directly or indirectly. For example, the most cited figures on the problem of evil are almost all from WEIRD societies (Western, Educated, Industrialized, Rich, Democratic).³⁸ Could this be shaping the debate in globally and historically peculiar ways largely invisible to most of its participants?

And finally, experimental philosophers of religion should also pay attention to how easily folk beliefs in the above areas are formed, given up, or revised. For example, imagine that Maria reports that the distribution of pain and pleasure in this world is epistemically more probable on Draper's hypothesis of indifference than on theism; however, let's also imagine that it takes almost no effort at all to get Maria to think differently (either changing her mind entirely or becoming agnostic on this issue). In a different vein, imagine that Tomás, a theist, reports that there is

³⁶ For example, Augustine in the ancient period and Marilyn Adams more recently.

³⁷ The defeat of evil is compared to the way in which ugly patches of color might enhance the beauty of the overall painting. It is not that the ugly patch is a flaw that is merely compensated for by the other parts of the painting; rather, the beauty of the whole is unachievable apart from the ugliness of that particular patch.

³⁸ See Henrich 2020 for an overview of how WEIRDness can skew our understanding of an issue.

plausibly some greater good that comes about because of the suffering of Rowe's fawn; however, when placed under cognitive-load (i.e., forced to think fast or has his reflective-reasoning weakened by other tasks) he denies that this is plausible. How would these results change our analysis of their initial reports? Between this sort of evidence (i.e., how easy it is to form or revise a belief), cross-cultural evidence, and the other sorts of evidence mentioned above, experimental philosophers of religion might get a very different picture of what the folk think about these issues.³⁹

These represent just a handful of the directions empirical inquiry into the problem of evil might take.

Conclusion

In this paper, we have undertaken an examination of the problem of evil through the lens of experimental philosophy, an approach that has demonstrated considerable promise in addressing some of philosophy's most enduring and complex issues. At the core of our investigation lies the phenomenon of suffering—a universal human experience that nevertheless elicits a remarkably diverse range of responses, being a central motivation for atheism in some and inspiring deeper faith in others. Further research applying the tools and resources of experimental philosophy to the problem of evil shows potential for uncovering the underlying cultural, psychological, and sociological factors that contribute to this diversity, thereby enriching our understanding of the problem of evil and its various philosophical formulations.

To conclude, we are calling for a rigorous, empirically informed approach to the problem of evil within the domain of experimental philosophy. By integrating empirical research with philosophical inquiry, we can achieve a more nuanced and comprehensive understanding of the complexities surrounding evil and suffering. This approach not only contributes to the philosophical discourse on the problem of evil but also exemplifies the broader potential of experimental philosophy to illuminate and challenge traditional philosophical problems.⁴⁰

References

- ADAMS, M. M. (1993) 'Aesthetic Goodness as a Solution to the Problem of Evil', in Arvind Sharma (ed.) *God, Truth and Reality: Essays in Honor of John Hick* (New York: St. Martin's), 46-61.
- (1999) *Horrendous Evils and the Goodness of God* (Ithaca: Cornell University Press).
- (2013) 'Ignorance, Instrumentality, Compensation, and the Problem of Evil', *Sophia*, 52, 7–26.
- ALEXANDER, D. AND D. JOHNSON (eds.) (2016) *Calvinism and the Problem of Evil* (Portland: Pickwick).
- ALSTON, W. P. (1991) 'The Inductive Argument From Evil and the Human Cognitive Condition', *Philosophical Perspectives*, 5, 29–67.
- AUGUSTINE, ST. (1964) *On Free Choice of the Will*, A. S. Benjamin and L. H. Hackstaff (trans.) (Indianapolis: The Bobbs-Merrill).
- BARRETT, J. L., and I. M. CHURCH (2013) 'Should CSR Give Atheists Epistemic Assurance? On Beer-Goggles, BFFs, and Skepticism Regarding Religious Beliefs', *The Monist*, 96, 311–24.
- BEIT-HALLAHMI, B. (2006) 'Atheists: A Psychological Profile', in M. Martin (ed.) *The Cambridge Companion to Atheism*, (Cambridge: Cambridge University Press), 300–318.

³⁹ We're grateful to Justin Barrett for pointing this out to us.

⁴⁰ This research was made possible thanks to the generous support of the John Templeton Foundation (grant ID: 61886). We're enormously grateful to Justin Barrett for providing feedback on this article for us.

- BERGMANN, M. (2001) 'Skeptical Theism and Rowe's New Evidential Argument From Evil', *Noûs* 35, 278–296.
- BRADSHAW, M., C. G. ELLISON, AND J. P. MARCUM (2010) 'Attachment to God, Images of God, and Psychological Distress in a Nationwide Sample of Presbyterians', *The International Journal for the Psychology of Religion*, 20, 130–47.
- CHISHOLM, R. (1968) 'The Defeat of Good and Evil', *Proceedings and Addresses of the American Philosophical Association*, 42, 21–38.
- CHURCH, I. M. (Forthcoming) 'Data Over Dogma: A Brief Introduction to Experimental Philosophy of Religion', *Philosophy Compass*.
- (2023) 'Experimental Philosophy of Religion: The Problem of Evil', in A. M. Bauer and S. Kornmesser (eds.) *The Compact Compendium of Experimental Philosophy* (Berlin: De Gruyter).
- CHURCH, I. M., R. CARLSON, AND J. BARRETT (2020) 'Evil Intuitions? The Problem of Evil, Experimental Philosophy, and the Need for Psychological Research', *Journal of Psychology and Theology*, 49, 126-141.
- CHURCH, I. M., BLAKE MCALLISTER, J. S. SPIEGEL, AND J. L. BARRETT (in progress) "Testing Folk Perceptions of Draper's Probabilistic Problem of Evil."
- CHURCH, I. M., I. WORCOL, AND J. 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*, 6, 4-19.
- CRAGUN, R. T. (2016) 'Nonreligion and Atheism' in D. Yamane (ed) *Handbook of Religion and Society* (New York: Springer), 301–320.
- DAVIES, B. (2006) *The Reality of God and the Problem of Evil* (London: Bloomsbury).
- DAVIS, S. T. (ed.). (2001) *Encountering Evil: Live Options in Theodicy*, 2nd edition (Louisville: Westminster John Knox).
- DE CRUZ, H. (2015a) 'Divine Hiddenness and the Cognitive Science of Religion', in A. Green and E. Stump (eds.) *Hidden Divinity and Religious Belief: New Perspectives* (Cambridge, UK: Cambridge University Press), 53–68.
- (2015b) 'Where Philosophical Intuitions Come From', *Australasian Journal of Philosophy*, 93, 233–249.
- (2017) 'Religious Disagreement: An Empirical Study Among Academic Philosophers', *Episteme*, 14, 71-87.
- (2018) 'Religious Beliefs and Philosophical Views: A Qualitative Study', *Res Philosophica*, 95, 477–504.
- DE CRUZ, H. AND R. NICHOLS (2016) *Advances in Religion, Cognitive Science, and Experimental Philosophy* (London: Bloomsbury).
- DOUGHERTY, T. (2014) 'Skeptical Theism', in E. N. Zalta and Uri Nodelman (eds.) *The Stanford Encyclopedia of Philosophy* (Stanford: Stanford University Press).
- DRAPER, P. (1989) 'Pain and Pleasure: An Evidential Problem for Theists', *Noûs*, 23, 331–50.
- DRAPER, P. AND RYAN NICHOLS (2013) 'Diagnosing Bias in Philosophy of Religion', *The Monist*, 96, 420–46.
- EKSTROM, L. W. (2021) *God, Suffering, and the Value of Free Will* (Oxford: Oxford University Press).
- FALES, E. (2013) 'Theodicy in a Vale of Tears', in J. P. McBrayer and D. Howard-Snyder (eds.) *The Blackwell Companion to the Problem of Evil* (Hoboken NJ: Wiley-Blackwell), 349–362.
- GOLDBERG, L. R. (1993) 'The Structure of Phenotypic Personality Traits', *American Psychologist*, 48, 26–34.

- Haidt, J. (2009) 'Moral Psychology and the Misunderstanding of Religion', in J. Schloss and M. Murray (eds.) *The Believing Primate : Scientific, Philosophical, and Theological Reflections on the Origin of Religion* (Oxford, UK: Oxford University Press), 278-291.
- HASKER, W. (2008) *The Triumph of God over Evil: Theodicy for a World of Suffering* (Downers Grove: InterVarsity Press).
- HENRICH, J. (2020) *The WEIRDest People in the World: How the West Became Psychologically Peculiar and Particularly Prosperous* (New York: Farrar, Straus and Giroux).
- HOWARD-SNYDER, D. (2009) 'Epistemic Humility, Arguments From Evil, and Moral Skepticism' in J. L. Kvanvig (ed.) *Oxford Studies in Philosophy of Religion*, Vol. 2 (Oxford: Oxford University Press), 17–57.
- JOHNSON, K. A., M. A. OKUN, A. B. COHEN, C. A. SHARP, AND J. N. HOOK (2019) 'Development and Validation of the Five-Factor LAMBI Measure of God Representations', *Psychology of Religion and Spirituality*, 11, 339–349.
- LEE, K. AND M. C. ASHTON (2004) 'Psychometric Properties of the HEXACO Personality Inventory', *Multivariate Behavioral Research*, 39, 329-358.
- LIM, D. F. (2017) 'Experimental Philosophy and Philosophy of Religion', *European Journal for Philosophy of Religion*, 9, 139–158.
- MACKIE, J. L. (1955) 'Evil and Omnipotence', *Mind*, 64, 200–212.
- MCALLISTER, B. (2020) 'The Perspectival Problem of Evil', *Faith and Philosophy*, 37, 421–450.
- MCALLISTER, B., I. M. CHURCH, P. REZKALLA, AND L. NGUYEN (2024) 'Empirical Challenges to the Evidential Problem of Evil', in J. Knobe and S. Nichols (eds.) *Oxford Studies in Experimental Philosophy*, Vol. 5 (Oxford: Oxford University Press).
- MURPHY, M. C. (2017) *God's Own Ethics: Norms of Divine Agency and the Argument From Evil* (New York: Oxford University Press).
- PLANTINGA, A. (1974) *God, Freedom, and Evil* (Grand Rapids, MI: Eerdmans).
- (2004) 'Supralapsarianism, or "O Felix Culpa"', in P. van Inwagen (ed.) *Christian Faith and the Problem of Evil* (Grand Rapids: Eerdmans), 1-25.
- REICHENBACH, B. 'Natural Evils and Natural Laws', *International Philosophical Quarterly*, 16, 179–196.
- ROWE, W. L. (1979) 'The Problem of Evil and Some Varieties of Atheism', *American Philosophical Quarterly*, 16, 335–341.
- (1996) 'The Evidential Argument From Evil: A Second Look', in D. Howard-Snyder (ed.) *The Evidential Argument From Evil* (Bloomington, IN: Indiana University Press), 262–285.
- SPEAK, D. (2015) *The Problem of Evil* (Malden, MA: Polity).
- SPIEGEL, J. S. (2011) 'On Free Will and Soul-Making: Complementary Approaches to the Problem of Evil', *Philosophia Christi*, Series II, 13, 405-413.
- STUMP, E. (1985) 'The Problem of Evil', *Faith and Philosophy*, 2, 392–423.
- (2010) *Wandering in Darkness: Narrative and the Problem of Suffering* (Oxford: Oxford University Press).
- (2011) 'Reply to Draper', in M. Bergmann, M. J. Murrar, and M. C. Rea (eds.) *Divine Evil?: The Moral Character of the God of Abraham* (Oxford: Oxford University Press).
- SWINBURNE, R. (1998) *Providence and the Problem of Evil* (Oxford: Oxford University Press).
- VAN INWAGEN, P. (2006) *The Problem of Evil: The Gifford Lectures Delivered in the University of St. Andrews in 2003* (New York: Oxford University Press).
- VAN LEEUWEN, Neil (2014) 'Religious Credence Is Not Factual Belief', *Cognition*, 133, 698–715.

- WYKSTRA, S. J. (1984) 'The Humean Obstacle to Evidential Arguments From Suffering: On Avoiding the Evils of "Appearance"', *International Journal for Philosophy of Religion*, 16, 73–93.
- (1996) 'Rowe's Noseeum Arguments from Evil', in D. Howard-Snyder (ed.) *The Evidential Argument From Evil* (Bloomington, IN: Indiana University Press), 126-150.