God and the Bayesian Conception of Evidence

David Manley

forthcoming in *Religious Studies*

Contemporary arguments for and against the existence of God are often formulated within a broadly Bayesian framework. Arguments of this sort focus on a specific feature of the world that is taken to provide probabilistic evidence for or against the existence of God: the existence of life in a 'fine-tuned' universe, the magnitude of suffering, divine hiddenness, etc. In each case, the idea is that things were more likely to be this way if God existed than if God did not exist—or the other way around.

Less attention, however, has been paid to the deeper question of *what it takes* for something to count as evidence for or against the existence of God. What exactly is being claimed when it is said that some feature of the world is *more or less likely* given the existence of God, and how should we go about assessing such a claim? This paper is about epistemological issues—and in particular, certain potential cognitive errors—that arise when we reason probabilistically about the existence of God. The moral is not that we should refrain from reasoning in this way, but that we should be mindful of potential errors when we do.

1. Introduction

The Bayesian conception of evidence is diachronic—or, at least, dyadic. What makes something a piece of evidence is the *difference it makes* between two information states. In the simplest case, a subject learns something and changes her beliefs by properly updating. We can then contrast her prior belief state with the belief state altered only by updating on what she learned. We then note the changes that propagate through her various degrees of belief (or "credences"). What the subject learned constitutes evidence for a hypothesis H just in case updating increases her credence in H. (And what she learned constitutes evidence against H just in case updating decreases her credence in H.) This, at least, is what we I will call the *subjective diachronic* conception of evidence.

On this approach, the notion of evidence is cashed out in terms of notion of properly updating—that is, *conditionalizing*—on one's prior credences. This means that, upon observing O, one's new credence in every H is equal to one's prior conditional credence in H given O.[1] And this conditional credence in turn is usefully thought of as a function of how (subjectively) probable O was given H,

[1] I am setting aside cases of Jeffrey conditionalization for simplicity, though the account of evidence can easily be extended to them (see Jeffrey 1983).

how probable O was given not H, and how probable H was *simpliciter*.[2] In particular, if O was more probable given H than given not-H, then the probability of H will increase: so O constitutes evidence for H.[3]

But things are not always this simple. First, there are cases where a subject *hadn't even considered* an important hypothesis before learning some new fact. Maybe the hypothesis occurred to her just as she learned the new fact, or just after. Still, intuitively, the new fact could still constitute evidence for that hypothesis. In such a case, we can't offer the simple before-and-after story given above about what makes something evidence for the subject. This is the *problem of old evidence*— or, equivalently, the problem of new hypotheses.[4] In a different kind of case, the subject *had* considered the relevant hypothesis, but had not considered the possibility of this particular kind of *evidence*. As a result, she had not considered in advance the conditional probability of this evidence given her various hypotheses. But without these prior conditional credences, she cannot follow the rule for updating on evidence.

In either kind of case, the subject lacks priors that are necessary for updating on evidence. And the very same issue arises for most of the evidence adduced in arguments for the existence of God. We are already aware of the existence of conscious life before asking ourselves how probable it is given the existence or non-existence of God. And presumably most of us are already aware of apparently pointless suffering before asking ourselves how probable such suffering is given the existence or non-existence of God. Likewise, at least for many people, when it comes to divine hiddenness, the vast emptiness of the universe, and so on.

Of course, this is not a special problem for purported evidence for the existence of God. Much of the evidence supporting our most important theories about the world, from physics to biology to psychology, consists of facts that we were aware of prior to formulating those theories: it is therefore old evidence. And we can hardly deny that many of these theories were first formulated precisely to fit the existing data. For a Bayesian account to explain how evidence works in these case, it must therefore allow for some alternative to the idea of temporally prior credences, at least for cases like these.

In what follows I will assume that some alternative notion of priors can be used to account for our old evidence in cosmology and biology, and likewise for purported evidence about the existence of God. The purpose of this paper is to

---

[2] Where 'p_{new}' and 'p' denote my new and old credence functions respectively, the updating rule is $p_{new}(H) = p(H|O)$; and the latter is defined as $p(H\&O)/p(O)$, at least where $p(O) > 0$. And this in turn is equivalent to $p(O|H)p(H) \,/\, [p(O|H)p(H) + p(O|{\sim}H)p({\sim}H)]$: this is 'Bayes' Theorem'. This is easy to grasp by thinking of the prior credence function as an region, where each hypothesis occupies a proportion of that region equal to its credence value. (For example, if H and O each have a prior credence of .5 and are independent, there are four equal subregions representing H&O, H&~O, ~H&O, ~H&~O.) The rule tells me that when I learn O, I simply *zoom in* on the area occupied by O; the proportion of the new area occupied by each hypothesis corresponds to its posterior credence.

[3] I am assuming the subject's credences form a (coherent) probability distribution.

[4] For some discussions of the problem, see Earman 1992, ch. 5; Glymour 1980, ch. 3; Howson and Urbach 1989, 272-75; Joyce, ch. 6.

examine some epistemological problems that result when we employ a substitute for chronologically prior credences. Since one cannot in such a case simply *remember* one's own prior credences, one must somehow *discern* or *reconstruct* an adequate substitute, while attempting to disregard one's knowledge of the evidence. But this procedure will be rife with subconscious biases that are well-known to cognitive psychology. And the problem will be especially dramatic in contexts where: (i) the alleged evidence is a pervasive part of our experience—such as the existence of conscious life, or of suffering; and (ii) we are considering hypotheses that are capable of explaining a wide range of possible outcomes—such as the hypothesis that an all-powerful mind exists.

2. Prior obligations

Before examining these potential errors in detail, let me begin with a few key points about prior probabilities.

i. Counterfactual priors? Faced with a case of absent priors, it is tempting to substitute some kind of counterfactual credence for a chronologically prior credence.[5] We can ask ourselves, for example, how confident *would I have been* in the hypothesis, had I considered it before learning the evidence? But counterfactuals like this are predictably fraught with problems: perhaps in the closest world where I considered the hypothesis before learning the evidence, I did so because I was slightly mad, or because I had different and misleading evidence. Such cases illustrate that I am not necessarily interested in the closest world where I considered the hypothesis and lack the evidence. I am interested only in scenarios where my standards for evaluating hypotheses are unchanged and I have no other evidence to go on.

Dealing with abstract priors, then, is not just a matter of trying to assess a simple counterfactual. It must somehow involve assign probabilities using one's actual epistemic norms within the context of a selectively impoverished information state—one that brackets one's knowledge of the evidence at issue.[6] Let's call the assignments that result from this procedure our *hypothetical priors*.

ii. Subjective vs. objective priors. Some theorists hold that there is an *objectively correct* assignment to the probability of, for example, animal suffering conditional on the existence of God. Then one can think of the problem of reconstructing absent priors as merely an instance of the more general problem of trying to match our credences to the objectively correct probabilities. On such a view, in addition to the subjective diachronic conception of evidence we have sketched, on which

---

[5] I am framing things as though the diachronic notion of evidence is the paradigmatic and solutions to the problem of old evidence will somehow be derivative. In fact, for reasons unrelated to the present paper, I think diachronic conditionalization can get things wrong and is compelling only because it coincides in the vast majority of cases with a better rule that is synchronic and can handle *de se* as well as *de dicto* evidence. (For an example where diachronic conditionalization gets things wrong, see the Shangri La case in Arntzenius 2003; for some proposals for a better rule, see Meacham 2008 and [Redacted].)

[6] See e.g. the discussions in Lange 1999, Eells and Fitelson 2000, and Meacham 2008.

$$\text{O is evidence for H iff } p_{\text{old}}(\text{O}\,|\,\text{H}) > p_{\text{old}}(\text{O}\,|\,\sim\text{H}),$$

there will also be an objective conception where '$p_{\text{objective}}$' replaces 'p_{old}'. Thus, something might count as 'evidence for H (for me)' in the subjective sense while not objectively being evidence for H at all. In this paper, however, I will remain neutral on the question of objective probabilities. Even if the epistemic norm of our subjective credences is to match the objective probabilities, the process of trying to do so in a case of absent priors will likewise involve an attempt to bracket some of one's evidence. And that is all that matters for my purposes.[7]

iii. Unconditional priors. For O to count as evidence for H, we have said, is for our priors to treat O as more probable given H than given not-H. So the general shape of an argument to the effect that some fact F is evidence for the existence of God, where 'p*' represents our hypothetical prior probability function, will be:

$$p^*(\text{O}\,|\,\textit{God exists}) \;>\; p^*(\text{O}\,|\,\textit{God doesn't exist})$$

And likewise, with the inequality reversed, if we want to make the case that O is evidence against the existence of God.

Importantly, this means that we do not even need to evaluate the unconditional prior probability of H in order to identify what counts as evidence for H. And this allows those who offer probabilistic arguments about God to sidestep the contentious issue of how one might go about assigning principled unconditional priors to the proposition *God exists.*[8] The principle above tells us we needn't fix an assignment for that proposition at all in order to ask ourselves what counts as evidence for or against the existence of God.

Despite this, we shouldn't lose sight of the fact that even very strong evidence can fail to move a hypothesis from the status of *very improbable* to the status of *more probable than not*. Suppose Jones, a perfectly ordinary person, wins the lottery. The fact that Jones won is strong evidence that the lottery was rigged for Jones, in the sense that it should significantly increase our credence that the lottery was rigged for Jones:

$$p_{\text{old}}(\textit{Jones wins}\,|\,\textit{Rigged for Jones}) \;\gg\; p_{\text{old}}(\textit{Jones wins}\,|\,\textit{Chance})$$

---

[7] In fact, in an objectivist framework it is easier to tell a story about what is going *wrong* if one falls prey to a cognitive bias in a case of absent priors (see section 3): the bias will tend to make one's assessment of the relevant conditional priors *objectively inaccurate*.

[8] Swinburne (1991) and others have argued that p^*(God) should be fairly high due to God's metaphysical simplicity. There is a problem with the connection between metaphysical simplicity is tied to theoretical simplicity: consider, for example, the hypothesis that all of my experiences are being fed to me by an evil demon who for whatever reason wants my experiences to be exactly like this. In explaining my experiences, this hypothesis trades the ontological complexity of realism about the external world for a highly complex and unlikely set of inner states in the evil demon. Positing in addition that the demon is metaphysically simple does not seem to help at all.

Of course we do not now conclude that the lottery was rigged for Jones, because the probability of *Rigged for Jones* was extremely low to begin with.[9] Indeed, if we had no more reason to think that Jones would rig the election than anyone else, his winning provides no evidence that the election is rigged in general: the probability of *Rigged for Jones* started off much lower than the probability of *Rigged*, and now they are roughly equal.

As with *Rigged* and *Rigged for Jones*, there is often a tradeoff between (on the one hand) a general hypothesis that has higher initial plausibility but does not make the evidence very probable, and (on the other) a specific hypothesis that has lower initial plausibility but makes the evidence more probable. And it is very important to keep track of which hypothesis is at issue when we are considering evidence. So, for example, suppose we are considering the general hypotheses *God exists* and *Only a material universe exists*. When considering how well these hypotheses predict various observations, it can be tempting to proceed as though we were considering much more specific hypotheses, like *God-construed-as-having-various-specific-goals exists*, or *Only a material universe-with-such-and-such-attributes exists*. The relevant evidence may be more probable conditional on these hypotheses, but this cost may well be offset by the lower initial plausibility of the more specific hypothesis.

3. Ad hocery and biases

We have seen that the Bayesian conception of evidence needs something like the notion of hypothetical priors in order to handle various examples from the sciences in which important theories are supported primarily by old evidence. Still, we often find it a bit suspect when a theory is formulated or modified precisely to suit the evidence. And our theory of evidence should be able to explain this reaction as well. If the process of forming hypothetical priors were completely unproblematic, it would be hard to say why it is better to specify a hypothesis and its predictions *before* gathering one's data—rather than mining the data after the fact to identify a hypothesis that it supports.

In short—what, if anything, is wrong with *ad hocery*? Consider two paradigmatic examples.

i. The layered moon. Lodovico delle Colombe famously held that the moon was perfectly smooth and spherical, and resisted Galileo's evidence to the contrary:

> [He] attempted to reconcile the old doctrine with the new observations, by asserting, that every part of the moon, which to the terrestrial observer appeared hollow and sunken, was in fact entirely and exactly filled up with a clear crystal substance, perfectly imperceptible by the senses, which restored to the moon her accurately spherical and smooth surface. (Bethune 106)

Note that the evidence adduced by Galileo is just what we would expect if this new hypothesis were true. "But," we want to say, "the crystal-layer hypothesis is so *ad hoc*!" And what's wrong with that, exactly? The theory of natural selection and the Big Bang theory are both 'post hoc' in the sense of being *developed after*

[9] Of course, if Jones's cousin created the system that picks the lottery winner, we might start to get worried. (President Mugabe famously won a lottery in Zimbabwe that was organized by a partially state-owned bank.)

encountering the evidence, and even 'ad hoc' in the sense of being *crafted specifically to explain* the evidence.

Galileo's own response was to reject the new hypothesis as absurd.[10] In Bayesian terms, this can be cashed out as the claim that it deserves a low prior probability. After all, there are as many hypotheses about layers of the moon as the are coherent combinations of ways for each layer to be. (For example, each layer might be smooth or irregular, and it might be opaque or translucent or transparent. If we limit ourselves just to two layers with only these features, we get two dozen competing hypotheses, a dozen of which involve a smooth outer surface.)[11] Faced in advance with all these hypotheses, even della Colombe would presumably not have assigned a very high credence to the specific hypothesis that the underlying surface is irregular and opaque, while the covering surface is smooth and transparent. And as a result, he probably would have assigned a much greater credence to (*Looks irregular* | *Is irregular*) than to (*Looks irregular* | *Is smooth*). And then the observation that it looks irregular would have been strong evidence that it is irregular—just as Galileo argued.

Since della Colombe tailored his hypothesis to Galileo's evidence, however, he probably did not even *consider* these other options. Instead, just one of them was salient and seemed pretty plausible—the one that conveniently explained the evidence at hand. As a result, della Colombe treated this hypothesis as though it were fairly probable conditional on the moon's being smooth, and he could dismiss Galileo's observations as providing little contrary evidence.

ii. Shy ESP. Whatever our prior credence in the reality of extrasensory perception, we surely get evidence against it when practitioners repeatedly fail to demonstrate their talents under experimental conditions. A common reply from the ESP camp is that there may be factors that make ESP fail to operate under experimental conditions. If *that* were true, we'd expect exactly the evidence we got! It is tempting to conclude that, since the experiments don't provide evidence against that hypothesis, they don't provide evidence against ESP!

But that would be fallacious. The experiments were designed to test the general hypothesis that ESP is real—*that* hypothesis on the whole does not predict the evidence. Of course, there is a *far more specific* hypothesis that does predict the new evidence—that ESP is real but shy of experiments. But that was a far less likely hypothesis to begin with. By ruling out the most plausible variants of the ESP hypothesis, the experiments do constitute strong evidence against the

---

[10] "Galileo met the argument in the manner most fitting, according to one of Aristotle's own maxims, that "it is foolish to refute absurd opinions with too much curiosity."" (Bethune 106)

[11] As Galileo mocked: "I am perfectly ready to believe [your theory] provided that with equal courtesy, I may be allowed to raise upon your smooth surface, crystal mountains (which nobody can perceive) ten times higher than those which I have actually seen and measured" (Bethune 106).

hypothesis that ESP is real— even though the evidence is just what we'd expect if a specific variant of that hypothesis were true.[12]

Again, if proponents of ESP were asked to assign probabilities in advance, they would likely not have treated the 'shy ESP' hypothesis as though it were nearly as probable as the general ESP hypothesis itself. There are all sorts of hypotheses about when ESP would work—that it would fail precisely under just those conditions used by experimenters to test it seems pretty implausible *ab initio*. But once it is *known* not to work under those conditions, that hypothesis seems more salient and therefore more plausible.

If I am right, there is nothing magical about assigning priors before encountering the relevant evidence, but it does fend off the tendency to decide that the outcome one happens to observe was likely all along, or even that it coheres especially well with one's favored hypothesis. On this way of thinking about it, the problem with ad-hoc hypotheses is purely psychological: they skew one's assessment of what should have been uncontaminated priors. And there is plenty of empirical data about mechanisms that underlie the phenomenon of probability judgments skewed by knowing the actual outcome.[13]

There are at least three (somewhat overlapping) areas of the literature on cognitive biases that are relevant:

1. *Hindsight bias* is the tendency for people's retroactive assignments of probability—as well as reports about what their own past probability assignments would have been—to be skewed by knowing the actual outcome. (This is sometimes called the 'Knew-it-all-along' effect.) This effect holds even when subjects are explicitly asked to make their judgments of probability as though they were unaware of the actual outcome (Hawkins & Hastie, 1990). The actual outcome can come to seem inevitable even if subjects are told it was a chance event: this has been dubbed 'creeping determinism'.

---

[12] When we rule out that ESP is both real and works in experimental conditions, we are obliged to renormalize and increase our credence *both* in the hypothesis that ESP is not real *and* in the hypothesis that ESP is real but only works in friendly conditions. But since only an ESP hypothesis was ruled out and renormalization involves increasing the probability of the surviving hypotheses *proportionally*, the experiments *do* disconfirm ESP.

[13] Suppose researchers test a drug on a group of subjects with liver cancer and find no statistically significant effect. They then analyze 20 subgroups, and find one for which the drug's effects are statistically significant, with a p-value of .05. (Note that we should *expect* one of 20 subgroups to yield statistical significance even if the drug has no effect.) In this scenario, regulatory agencies would be to require a further study aimed specifically at the new hypothesis in order to rescue it from 'ad hoc' status, regardless of its medical plausibility. (The subgroup might be 'patients named Bob', or 'patients with co-morbid cirrhosis'.) For a Bayesian, the problem with all this is the emphasis on a single arbitrary threshold of statistical significance. Realistic priors should differentiate the 'Bob' and 'cirrhosis' subgroups, requiring far more impressive results to confirm the former. On this view, whether a given hypothesis is *ad hoc* shouldn't matter in principle— though due to the kinds of biases discussed in the text, any *priors* assigned after the fact would be suspect. From a Bayesian standpoint, the current system uses the cost of running an additional study as a kind of proxy for high priors.

2. The *availability heuristic* is the general tendency for people to reason in a way that overemphasizes what comes easily to mind. Notably, "the plausibility of the scenarios that come to mind, or the [ease] of producing them, then serve as a clue to the likelihood of an event" ((Tversky and Kahnemann 1973, 229). This has been borne out in subsequent research: instructing subjects to imagine an event increases their probabilistic expectations for that event (Carroll 1978, Gregory et al. 1982), and people are far more influenced in their probability judgments by concrete examples than by statistics and base rates (Nisbett and Borgida 1975).

3. Salience may also increase estimates of probability by causing subjects to consider fewer *alternative outcomes*. Subjects who consider fewer possible outcomes have higher likelihood estimates than those who consider fewer (Dougherty et. al. 1997, Gregory et al. 1982)— and perhaps for that reason, those who generated outcomes on their own were less confident than those who had the outcomes suggested to them (Koehler 1994).

4. A well-known aspect of *confirmation bias* is the tendency to interpret new information in such a way that it supports hypotheses to which we are already inclined, whether epistemically or emotionally (Klayman 1994, Kohler 1993). Ambiguous evidence in particular is likely to be manipulated in this way: on a Bayesian model with hypothetical priors, it is natural to model this as overly generous credences for known events conditional on one's favored hypothesis. The effect of confirmation bias can be somewhat mitigated if we are forced to assign credences to potential outcomes *in advance* rather than allowing ourselves to find ways of making the outcomes fit our hypothesis.

There is sure to be a good deal of interdependence between these four phenomena, but the central point for our purposes is that they all raise concerns for the practice of generating hypothetical priors for some piece of already-known evidence. Hindsight bias will tend to make our priors more confident than they would otherwise be; the salience of the known outcome will tend to suppress the consideration of alternative possibilities, which also increases likelihood judgments; and if we happen to feel confident in (or emotionally attached to) a given hypothesis, we will tend towards generosity in estimating its ability to predict actual outcomes.

These combined difficulties are exacerbated further when trying to assign credences to outcomes given the existence of God. The hypothesis of theism itself —the existence of an all-powerful mind—is capable of explaining a staggering array of possible outcomes, many literally unthinkable by finite minds. But the sheer difficulty of considering those alternate possibilities will be exacerbated by the cognitive biases just mentioned. In addition, the evidence adduced in arguments for and against the existence of God tends to be a conspicuous part of our experience, such as the the existence of life or of apparently pointless suffering. So not only is it salient and available, but its sheer cognitive pervasiveness will further exacerbate the tendency not to envision alternatives.

There are, however, some mental exercises that can help combat cognitive biases (Larrick 2004). One is to force oneself to consider several possible alternative outcomes and attempt to imagine them concretely. Another is to 'consider the opposite' of the evidence that one has actually acquired, and ask oneself how one would have reacted if one had learned that instead. For example,

if we find O to be evidence for H, would we also have treated ~O as evidence for H (or even neutral)? Or, if we deny that O is evidence for against H, would we nevertheless have treated ~O as evidence for H? If so, something has gone badly wrong. For example, suppose someone treats the ESP experiments as providing no evidence against ESP, but would have treated the opposite outcome—the experiments indicate ESP is real!—as evidence *for* ESP. The result is incoherent. After all, since O and ~O are mutually exclusive and jointly exhaustive,

$$p^*(O \mid H) > p^*(O \mid {\sim}H)$$

will hold just in case

$$p^*({\sim}O \mid {\sim}H) > p^*({\sim}O \mid H).$$

So by our principle of evidence, O is evidence for H iff ~O is evidence for ~H. In other words, if a piece of evidence is good for the goose, its negation would have been bad for the goose!

4. God and the priors

With these conceptual tools in hand, let us turn now in detail to some examples of probabilistic arguments about the existence of God and the hypothetical priors needed to evaluate them.

i. The emergence of persons. The existence of rational agents in the world is arguably evidence for the existence of God. Theists frequently make the case that a divine person would be motivated to create persons with whom to enter into relationship. In contrast, various reasons are adduced why a physical universe may have been unlikely to produce any life—the improbability of the emergence of self-replication on a given planet, or even the improbability of the 'fine-tuned' constants that allow for the possibility of large particles or planets in the first place.

Here the cognitive biases we have considered may tempt the atheist to downplay the various ways in which a material universe might have evolved that do not give rise to life— including various difficult-to-imagine scenarios where the laws or constants differ significantly. The *actual* evolution of the universe—complete with large particles and planets—seems in retrospect particularly salient and likely. And so does the actual path of evolution, with self-replicating life arising by chance and conscious and rational beings emerging in due course.

To counteract the pull of this 'creeping determinism', it can be useful to try some mental exercises. First, actively dwelling on possible alternatives will help mitigate mere availability bias. For example, we can reflect on the various points at which key ancestors might have died, thus ending our evolutionary branch and arguably leaving the world without persons. Second, we can try the "goose" principle a useful check on any gut feeling that the existence of rational life provides absolutely no evidence of the existence of God. Would the lack of any rational life in the universe—viewed, of course, from the abstract perspective of hypothetical priors—provide any evidence *against* the existence of God? (Compare arguments to the effect that vast cold regions of space completely devoid of life provide evidence against the existence of God.) If one's gut responds positively to *that* question, it is worth reassessing whether the existence of rational life provides evidence for God—on pain of incoherence.

However, given the cognitive biases that are likely to be at work in constructing hypothetical priors about the existence of rational life, it would be foolhardy to give much weight to our untutored intuitions. But the alternative is extremely arduous. Assigning a principled probability to the emergence of life requires answering very speculative questions on the fringes of scientific inquiry—questions about whether we should think of the parameters of our universe as finely tuned and therefore unlikely to have given rise to life; questions about the likely number of universes in the multiverse (if there is one), about the number of planets in the universe, about how likely a given planet is to give rise to life, etc. Mere intuitions should have little force when constructing these hypothetical priors on this very vexed and multifaceted question.

ii. Apparently pointless suffering. Consider this exchange:

> A: Surely the magnitude and apparently capricious distribution of suffering is evidence against the existence of God.
>
> B: But supposing God has good reasons to allow this very magnitude and distribution of suffering, then we would expect to experience exactly what we do experience!

What B says is perfectly true, but it would be a mistake to conclude from this that suffering provides *no* evidence against the existence of God. Everything hinges, of course, on the probability of this more specific theistic hypothesis given theism, which in turn will influence the probability of the evidence given theism. The epistemological difficulty lies with trying to come up with an impartial assessment of these probabilities. Collectively, the data from cognitive science adduced above suggest that the attempt to assign these probabilities retroactively is likely to be subject to various cognitive biases.

First, our imaginations are limited and it is extremely difficult to assess a 'full range' of possibilities involving God having reasons to allow various degrees and distributions of suffering. This problem is especially likely to be acute in this case because suffering is so much a part of our lives that it is hard to imagine worlds without it, and also because it can seem futile even to guess at the options open to an omnipotent deity. In addition, the very pervasiveness of suffering can make it seem particularly inevitable. We are left with a single very salient possibility— the one consistent with how things turned out— and a bunch of other possibilities that are neither salient nor even well-defined. This is, of course, not the situation we would be in, hypothetically, if we were somehow impartial bystanders to the creation of the universe with no knowledge of how much suffering, if any, there would be.

Some deny that apparently pointless suffering provides any evidence against the existence of God. This might be for two reasons—either they do not take themselves as capable of assigning any conditional credences of the relevant kind, or they do assign such credences but treat p^*(*Apparently pointless suffering* | *God exists*) to be no lower than p^*(*Apparently pointless suffering* | *God does not exist*). In the latter case, a useful hedge against the potential role of confirmation bias is to ask oneself whether one would count it as evidence for the existence of God if one discovered that, against all odds, all suffering has a clear point to it. If so—the goose principle entails that these two conditional priors cannot be equal.

Those who refuse to assign credences about suffering given theism typically do so out of an aversion to divine psychologizing. Converging with strains of apophaticism in Judeo-Christian traditions, 'skeptical theists' have argued that suffering (in its actual magnitude and distribution) provides no evidence against theism, because we are not in a position to assign a value to the probability of that suffering given God.[14] On this view, we simply do not know what to expect from a divinely created world when it comes to the magnitude and distribution of evil. For some, this is because we should not trust our modal intuitions at all when it comes to such out-of-the-ordinary things as counterfactuals about divine action; others call for modesty in particular about our insight into the workings of the divine mind.

Again, it is worth asking how such theorists would have responded to the discovery that the world actually contains no pointless suffering at all. If that would have counted as evidence *for* the existence of God, there are at least *imprecise* values for the relevant conditional credences at work—and perhaps the aversion to divine psychologizing is not absolute.

iii. The problem of proportion. One of the facts sometimes taken to constitute at least *some* evidence against the existence of God is the sheer lack of proportion in the universe, if its primary point was to create rational life. For example, the vast majority of hundreds of millions of species ever to have lived have come and gone, crushed in a tremendous avalanche of apparent evolutionary drift. What's alleged to be unlikely given theism is the "tremendous wastefulness of it, the tremendous cruelty of it, the tremendous caprice of it, the tremendous tinkering and incompetence of it".[15]

On the contrary, comes the reply: "the theory of biological evolution is simply irrelevant to the truth of Christian theism," because theism is perfectly consistent with slow and meandering evolution.[16] After all, being eternal, it's not as though God was in a hurry. Indeed there is a *version* of theism, one on which God prefers to create slowly and deliberately, and on which we would expect exactly something like evolution to be the case. Evolution is no evidence at all against such a version of theism. Therefore evolution is no evidence at all against theism.

The problem with this reply should be apparent. Everything hinges on the prior probability of this version of theism— and we may not be able to assign a value impartially when we already know that evolution is true.

> B: Well, I find it quite plausible that God would create slowly. There's no special reason to think God would create instantaneously.
>
> A: Actually, I suspect even your the hypothetical probability that God would create like this can't be very high on reflection. After all, if paleontologists had discovered that all living species popped into existence simultaneously and fully formed, *surely* that would have been evidence for the existence of God. But it

---

[14] e.g. Bergmann 2009; see also also Van Inwagen 1995, Wykstra 1984, Bergman 2001.

[15] The quote is from Christopher Hitchens in a debate with William Lane Craig (Craig & Hitchens 2009).

[16] Craig in (Craig & Hitchens, 2009).

> would *not* have been evidence for the existence of God if in advance we were pretty sure God would create slowly and deliberately. If discovering the sudden emergence of species would supports theism, then discovering gradual emergence of species supports atheism— by the goose principle.

A related point can be illustrated from this discussion between Tim Maudlin and Gary Gutting:

> TM: No one looking at the vast extent of the universe and the completely random location of homo sapiens within it (in both space and time) could seriously maintain that the whole thing was intentionally created for us…
>
> GG: I don't see why the extent of the universe and our nonprivileged spatio-temporal position within it says anything about whether we have some special role in the universe. (Gutting 2014)

Maudlin's remark suggests that he thinks our location in the universe is *extremely strong* evidence against theism, while Gutting's remark suggests he thinks it is *no evidence at all*. Gutting goes on:

> I think we need to distinguish different sorts of theism. There are versions of theism that, like a literal reading of Genesis, are inconsistent with what we know about cosmology. But there are also versions that don't require any specific story about the extent of the physical universe or our location in it.

Of course, it is almost trivially true that there are 'versions of theism' that are immune to cosmological evidence of the sort Maudlin adduces. The crucial question is about how that evidence bears on theism in general.

Suppose, for example, that what C.S. Lewis calls the 'discarded image', the cosmological picture of late antiquity and of the first millennium of Christendom, had been borne out by scientific inquiry (Lewis 1964). On this picture all the heavenly bodies revolve in perfect circles around the Earth, including that great black and star-studded sphere: the *stellatum*. The whole thing resembles a kind of mechanical Faberge Egg, with the cosmic significance of earthly life embodied in its very structure. Surely if this picture had turned out to be true, it would have constituted at least some evidence for the existence of a Designer, relative to our hypothetical priors. (Early theists widely took it to constitute such evidence, and they were not wrong about this!)

It is worth stressing that the goose principle is merely a check on the action of cognitive biases: it does not allow us to establish the *extent* of evidence the problem of proportion provides against the existence of God.[17] But it can be useful in helping us visualize the *alternative* to a vast, empty, and cold universe—and whether or not such an alternative fits better with atheism or theism. A *modest* evidential argument from proportion requires only that the probability of a vast, empty, and cold universe given theism is lower than the probability of a vast, empty, and cold universe given atheism.

---

[17] It is consistent with the goose principle that ~O would have provided extremely strong evidence for ~H, but O provides only very weak evidence for H. (For example, learning that my house is silent and dark is strong evidence that there is not an enormous party for clowns going on there; but learning that it is not silent and dark is only very weak evidence that there is a such a party going on.)

But perhaps Gutting would want to deny that the early apologists were right to take our apparently privileged position as evidence for theism:

> T.M.: Theism, as religious people typically hold it, does not merely state that some entity created the universe, but that the universe was created specifically with humans in mind as the most important part of creation. If we have any understanding at all of how an intelligent agent capable of creating the material universe would act if it had such an intention, we would say it would not create the huge structure we see, most of it completely irrelevant for life on Earth, with the Earth in such a seemingly random location, and with humans appearing only after a long and rather random course of evolution.
>
> G.G.: Maybe, but that conclusion doesn't follow from scientific cosmology; it's based on further assumptions about what a creator would want — and how the creator would go about achieving it.

The reluctance to assign desires and plans to God suggests another possible move, namely the analog of skeptical theism for cosmological evidence. On this view, we really have no good way to assign a value to $p^*(\textit{Vast empty universe} \mid \textit{God})$. But if this is so, Gutting would have to respond in the very same way to early Christian apologists using our privileged cosmic position as evidence for theism: that's no evidence at all because we have no way to know what kind of universe to expect from God. And it also precludes the use of anything like a Fine-Tuning Argument, since we would likewise have no way to assess things like $p^*(\textit{Fine Tuning} \mid \textit{God})$.

5. The upshot

Evidential arguments for and against the existence of God ask us to do something we are not very good at. They typically require us to set aside our knowledge of some very salient facts in order to reconstruct the hypothetical probability of those facts given competing hypotheses. There may be no alternative if we want to assess evidence for the existence of God, but this process is beset with epistemic danger and should be undertaken with vigilance.

Works Cited

Arntzenius, F. 2003, Self-locating Beliefs, Reflection, Conditionalization and Dutch Books", *Journal of Philosophy*, pp. 356-70

Bergmann, Michael, 2001, Skeptical Theism and Rowe's New Evidential Argument from Evil, *Noûs*, 35: 278–296.

Bergmann, Michael, 2009, Skeptical Theism and the Problem of Evil, in Flint, Thomas and Rea, Michael, *Oxford Handbook of Philosophical Theology*, Oxford: Oxford University Press, 375–399.

Bethune, John Elliot Drinkwater, 1830. *The life of Galileo Galilei, with illustrations of the advancement of experimental philosophy.* Boston: William Hyde & Co.

Carroll, J.S. 1978 The effect of imagining an event on expectations for the event: An interpretation in terms of the availability heuristic. *Journal of Experimental Social Psychology*, 14: 88–96

Craig, William Lane and Christopher Hitchens, Debate Transcript, Biola University, April 2009. Transcript accessed at: http://hitchensdebates.blogspot.com

Dougherty, Michael, Charles F. Gettys, and Rickey P. Thomas. 1997. The Role of Mental Simulation in Judgments of Likelihood. *Organizational Behavior and Human Decision Processes,* 70.2: 135-148.

Eells, Ellery, and Branden Fitelson, 2000, Measuring Confirmation and Evidence, Journal of Philosophy, 97: 663–672.

Gregory, W. Larry, Robert Cialdini, and Kathleen Carpenter 1982. Self-relevant scenarios as mediators of likelihood estimates and compliance: Does imagining make it so? *Journal of Personality and Social Psychology*, 43: 89–99

Gutting, Gary, 2014. Modern Cosmology Versus God's Creation (Interview with Tim Maudlin), *The Stone.* http://opinionator.blogs.nytimes.com/

Halvorson, Hans. Pre-print. A probability problem in the fine-tuning argument. Phil Sci Archive. http://philsci-archive.pitt.edu/11004/

Hawkins, S. A., & Hastie, R. 1990. Hindsight: Biased judgments of past events after the outcomes are known. *Psychological Bulletin*,107: 311-327.

Howson, Colin. *Objecting to God*. Cambridge: Cambridge University Press.

Koehler, J. J. 1993. The influence of prior beliefs on scientific judgments of evidence quality. *Organizational Behavior and Human Decision Processes*, 56,:28-55.

Koehler, Derek J. 1994. Hypothesis Generation and Confidence in Judgment. J*ournal of Experimental Psychology: Learning, Memory, and Cognition*. 20: 461-469.

Klayman, Joshua. 1995 Varieties of Confirmation Bias, *Psychology of Learning and Motivation, Academic Press*, 32: 385-418,

Lange, Marc, 1999, Calibration and the Epistemological Role of Bayesian Conditionalization, *Journal of Philosophy*, 96: 294–324.

Larrick, Richard, 2004. Debiasing. In: *The Blackwell Handbook of Judgment and Decision Making.* Malden, Mass: Blackwell.

Levi, Ariel S. and John B. Prior. 1987. Use of the availability heuristic in probability estimates of future events: The effects of imagining outcomes

versus imagining reasons. *Organizational Behavior and Human Decision Processes.* 40.2: 219-234

Lewis, C.S. 1964. *The Discarded Image*. Cambridge: Cambridge UP.

Meacham, Christopher J.G. 2008. Sleeping Beauty and the Dynamics of De Se Beliefs. *Philosophical Studies*, 138: 245-269.

Nisbett, Richard, and E. Borgida. 1975. Attribution and the psychology of prediction. *Journal of Personality and Social Psychology*, 32: 932-943.

Tversky, Amos and Daniel Kahneman 1973. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*. 5 (2): 207–232.

van Inwagen, Peter, 1995. *God, Knowledge, and Mystery.* Ithaca: Cornell University Press.

Wykstra, Stephen, 1984, The human obstacle to evidential arguments from suffering: on avoiding the evils of 'appearance,' *International Journal for Philosophy of Religion*, 16: 73–94.