Learning not to be Naïve: A comment on the exchange between Perrine/Wykstra and Draper[1]

Lara Buchak, UC Berkeley

ABSTRACT: Does postulating skeptical theism undermine the claim that evil strongly confirms atheism over theism? According to Perrine and Wykstra, it does undermine the claim, because evil is no more likely on atheism than on skeptical theism. According to Draper, it does not undermine the claim, because evil is much more likely on atheism than on theism in general. I show that the probability facts alone do not resolve their disagreement, which ultimately rests on which updating procedure – conditionalizing or updating on a conditional – fits both the evidence and how we ought to take that evidence into account.

Does postulating *skeptical theism* undermine the claim that the amount and type of evil in our world is evidence (strong or weak) against God's existence? Participants in this debate are ultimately interested in the relative probability of theism and atheism; whether evil confirms atheism over theism; and if it does, to what degree. A crucial issue is how the probabilities of these two hypotheses shift when we come to believe, through trying to construct a theodicy, that no satisfying positive account of why God permits the amount of evil in our world is forthcoming: that *naïve theism* is false and thus that skeptical theism is the only viable version of theism.[2] According to Perrine and Wykstra, the relevant comparison is between the best version of each theory, and comparing skeptical theism with (the best version of) atheism will show that evil does not strongly confirm atheism over theism. According to Draper, the relevant comparison is between atheism and theism full stop, and comparing these will show that evil does strongly confirm atheism over theism. Their disagreement is spelled out in a particular example of belief updating which they both discuss.[3] Examining how the probability facts change in this example will help us see what the disagreement ultimately rests on, and whether there is a way forward for the skeptical theist.

Here is the example. Consider two aliens, Natty (a naturalist) and Theo (a theist), who learn the empirical facts about our world in a particular order, while making some predictions on the basis of their theories N and T (I will use TS to stand for skeptical theism and T$\overline{S}$ to stand for naïve theism, making the simplifying assumption that these are the only two theistic options). They "first make predictions about pain and pleasure. Then, taking into account what they've learned about pain and pleasure, they make predictions about flourishing and languishing…Finally…they make predictions about triumph and tragedy" (Draper 5). According to Draper, Natty's predictions will be on the whole more accurate than

---

[1] Forthcoming in Skeptical Theism: New Essays, eds. Trent Dougherty and Justin McBrayer (OUP).

[2] What participants in this debate typically hold is that naïve theism is *very unlikely* to be true, given that we lack a positive account of why God allows evil. It will simplify the discussion to assume that naïve theism has been ruled out entirely, and that skeptical theism is the only alternative.

[3] Original example Draper. Here I follow Draper (2014: 5).

Theo's.[4] In probabilistic terms, if E_1, E_2, …, E_n are the data, and b is the background information, then $p(E_1 \mid N \,\&\, b) >> p(E_1 \mid T \,\&\, b)$, $p(E_2 \mid E_1 \,\&\, N \,\&\, b) >> p(E_2 \mid E_1 \,\&\, T \,\&\, b)$, and so on for successive data points (though Natty's predictions needn't all be *much* more accurate). If D, the total data of good and evil, is the conjunction of E_1, E_2, …, E_n, then $p(D \mid N \,\&\, b) >> p(D \mid T \,\&\, b)$. And as long as we don't have $p(T \mid b) >> p(N \mid b)$,[5] application of Bayes' Theorem shows that $p(T \mid D \,\&\, b) < ½$. None of this is controversial to participants in the debate if the only theistic possibility on the table is naïve theism.

What difference could skeptical theism make to this scenario? According to Perrine and Wykstra, Theo could develop as follows.[6] He begins as a naïve theist – or, more accurately, assigns most of theism's probability to naïve theism and only a tiny amount to alternative precisifications of theism – but then notices that his initial assumptions about what God values get him into trouble. As a result, "chastened by [his] failed predictions" and "given [his] commitment to theism" he now "shifts much of the probability he had assigned to naïve theism to a form of moderate skeptical theism."[7] Of course, once he has shifted most of theism's probability to skeptical theism, then his theism predicts the data just as well as Natty's naturalism.

What is Perrine and Wykstra's claim here, in probabilistic terms, and how is it supposed to undermine the above argument that $p(T \mid D \,\&\, b) < ½$? Draper mentions three interpretive possibilities, two of which are important for our purposes.[8] The first is that Theo's initial predictions remain the same and are inaccurate ($p(E_1 \mid N \,\&\, b) >> p(E_1 \mid T \,\&\, b)$), but thereafter Theo's theism becomes skeptical theism and his subsequent predictions are in line with those of naturalism ($p(E_2 \mid E_1 \,\&\, N \,\&\, b) = p(E_2 \mid E_1 \,\&\, T \,\&\, b)$ and so forth). As Draper also points out, however, this set of probabilities still implies that $p(D \mid N \,\&\, b) >> p(D \mid T \,\&\, b)$. The second interpretation, which Draper also argues against, is more interesting: upon realizing that naïve theism is not viable, Theo adopts skeptical theism,[9] and then predicts all of the data anew (including the initial data point). Since all participants in this debate agree that $p(D \mid TS \,\&\, b) = p(D \mid N \,\&\, b)$, then if $p(TS \mid T \,\&\, b) \approx 1$, it follows that $p(D \mid T \,\&\, b) \approx p(D \mid N \,\&\, b)$, and so D is not evidence for atheism over theism, strong evidence or otherwise.

[4] On p. [14], Perrine and Wykstra assign a particular probability distribution to Theo, though unfortunately p(T) is not given. One might assume from the particular values they assign that they mean to imply $p(T) = 0.5$. However, if, in parallel, Natty assigns $p(N) = 0.5$, this seems to imply that Theo and Natty assign the same probability as each other to N and T, obscuring the sense in which Theo is a theist (in Perrine and Wkystra's words, *commited* to theism) and Natty a naturalist. Therefore, I will assume that Theo assigns $p(T) = 1$ and Natty $p(N) = 1$, since this seems to better capture the assumptions that both sets of authors are making. Since what matters are the probabilities the *observer* assigns to each hypothesis at the end of the experiment, it won't make a difference to the debate. I will also assume for simplicity that T and N are mutually exclusive and exhaustive.
[5] See Draper (2014: [12]).
[6] Perrine/Wykstra (2014: [12-14]).
[7] Id., p. [14]. I will hereafter drop the modifier "moderate" from "moderate skeptical theism."
[8] Draper (2014: [18-19]).
[9] Or assigns he most of his probability to skeptical theism, as in Perrine and Wykstra's discussion: again, for mathematical simplicity I am assuming that alternatives to skeptical theism have been completely ruled out.

Draper suggests a problem for this version of Perrine and Wykstra's argument, by way of analogy. Suppose we have four urns, three of which (#1, #3, #4) contain many more yellow balls than purple balls, and one of which (#2) contains more purple balls than yellow balls. Balls are being drawn from one of the urns, and we are interested in how our probabilities about the urn they were drawn from should change in response to seeing the colors of these balls. Theory T12 says the balls are being drawn from urn #1 or urn #2, and comes in two versions: T1, which says they are drawn from urn #1, and T2, which says they are drawn from urn #2. Theory T34 says the balls are being drawn from urn #3 or urn #4. Upon seeing a yellow ball, two things ought to happen. First, a defender of T12 ought to shift probability from T2 to T1: on the assumption that T12 is correct, T1 is much more likely to be correct than T2. For the same reason, an impartial observer should raise her conditional probability p(T1 | T12). But second, and crucially, an impartial observer should *also* shift some of her probability from T12 to T34: as Draper puts it, the data favors T34 over T12.[10] And, according to Draper, these same two facts hold in the case of Theo and Natty. While Theo should raise his probability for skeptical theism (and we the observer should raise our conditional probability for skeptical theism given theism), we the observer should *also* lower our probability for theism overall: the data supports atheism over theism.

Draper's conclusion that the data supports T34 over T12 in the urn example is clearly correct. So what exact mistake is being made by the adherent of T12 who says "I now know that T1 is the best version of my theory, and the data does not support T34 over T1, so the data does not support T34 over T12"? (This is supposed to be analogous to Perrine's and Wykstra's claim that the data does not support atheism over theism *in general*, once we notice that skeptical theism is the best version of theism.) The move of letting the adherent make his view more specific in response to the data, and the thought that the question of which general theory the data supports is answered by looking at his prediction on the more specific theory, obscure an important point. When we learn that T2 is likely false, this has an effect not just on the relative probabilities of T1 and T2 conditional on the assumption that T12 is true, but also on

---

[10] As an example: let's say our antecedent probabilities are p(T1) = p(T2) = p(T3) = p(T4) = 0.25, and so p(T12) = p(T34) = 0.5 and p(T1 | T12) = p(T2 | T12) = 0.5. Let us also assume, following Draper, that p(Y | T1) = p(Y | T3) = p(Y | T4) = 0.9, and p(Y | T2) = 0.003. Then we have:

p(Y) = p(Y | T1)p(T1) + p(Y | T2)p(T2) + p(Y | T3)p(T3) + p(Y | T4)p(T4) = 0.67575.
p(T1 | Y) = p(Y | T1)p(T1)/p(Y) = 0.333 (and similarly for p(T3 | Y) and p(T4 | Y).
p(T2 | Y) = p(Y | T2)p(T2)/p(Y) = 0.001
p(Y | T12) = p(Y | T1)p(T1 | T12) + p(Y | T2)p(T2 | T12) = 0.4515
p(Y | T34) = p(Y | T3)p(T3 | T34) + p(Y | T4)p(T4 | T34) = 0.9.
p(T12 | Y) = p(Y | T12)p(T12)/p(Y) = 0.334 (another way to see this is that p(T12 | Y) = p(T1 v T2 | Y) = p(T1 | Y) + p(T2 | Y)
p(T34 | Y) = 0.666
p(Y & T12 | T12 & T1) = p(Y | T1)
p(T1 | Y & T12) = p(Y & T12 | T12 & T1)p(T1 | T12)/p(Y | T12) = 0.997

Upon learning Y, we conditionalize on Y, and so p_{new}(T1 | T12) = 0.997 and p_{new}(T12) = 0.334.
The particular values here are not crucial. As long as p(T2) > 0, we will have p_{new}(T1 | T12) > p(T1 | T12) and p_{new}(T12) < p(T12).

the relative unconditional probabilities of T12 and T34 – and it has an effect precisely by eliminating a previously viable possibility for one theory without doing the same for the other theory. (Technically, it assigns low probability to a previously higher-probability possibility, but let us assume for the sake of discussion that it eliminates the possibility altogether.) A better picture, according to Bayesianism, is the following (it assumes both that T12 and T34 are initially given equal probability by the observer, and that the disjuncts of T12 are initially given equal probability by the observer, but the analogy can be made more general). We begin with two representatives of T12, one who adheres to T1 and one who adheres to T2, and two representatives of T34 – so that the relative number of representatives is equal to the relative probability we assign to each theory. When the data comes in, the representative of T2 is eliminated. So while the spokesperson for T12 will now be our T1 representative, there will be more spokespeople for T34 overall. If the situation between Theo and Natty when they learn D is like the situation between the defender of T12 and the defender of T34 when they learn Y, then a better analogy for what happens in response to discovering that naïve theism isn't viable would mirror the analogy here, with the relative size of the voice in favor of theism decreasing.[11]

Therefore, if what we are interested in is how conditionalizing on D impacts the relative probability of theism and naturalism, then Draper is correct: even though conditionalizing on D should make theists become skeptical theists, D is evidence for atheism over theism.

Is there anything then to be said for Perrine's and Wykstra's position? Might there be other legitimate ways of updating on the data? Yes. We can model Natty's and Theo's situation in such a way that the evidence doesn't support atheism over theism – so that the observer ought not to lower her probability for theism. (Which model – this model or Draper's – is *accurate* will ultimately depend on certain further facts, as we will see.) The mathematical details aren't explicitly spelled out by Perrine and Wykstra, but this model makes good on what seems to be the primary point driving their argument: that what the data about evil supports is the *conditional claim* that if theism is true then skeptical theism is true.[12] The existence of such a model shows that, indeed, the question of whether the data of good and evil supports atheism even in light of the skeptical theist hypothesis turns not primarily on a mathematical question about what follows from the probability calculus and Bayesian updating, but on a more fundamental question about how to appropriately characterize the role of skeptical theism in the argument.

Let us step back and consider the general phenomenon of how probability shifts in response to evidence. A helpful way to think about this phenomenon is as follows, following van Fraassen (1989:

[11] A similar point holds if we are meant to interpret Theo as assigning $p(T) = 0.5$ rather than $p(T) = 1$ (see footnote 3): as he rules out non-skeptical theism, $p(TS \mid T)$ increases, but $p(T)$ decreases.
[12] See, for example, claim "D" about Granularism on p. 9.

161-2). Think of the entire space of hypotheses as a Venn diagram, where each region in the diagram specifies the truth-value of all the propositions we care about. We have a unit's worth of "mud" to distribute across the space, and the amount of mud in each region is the probability assigned to the proposition describing that region. So, for example, in the diagrams below (Figure 1), the amount of mud in the shaded region represents the probability of $H\bar{E}$ and of $A\bar{B}$, respectively; the amount of mud surrounded by the thick line represents the probability of H and of A, respectively. If we think of unconditional probability as a primitive, and define conditional probability in terms of it,[13] then conditional probability can be read off the diagram as well: the conditional probability of $\bar{E}$ given H is the ratio of the amount of mud in the $H\bar{E}$ region to the total amount of mud in the H region (the proportion of the thick-lined region that is shaded).

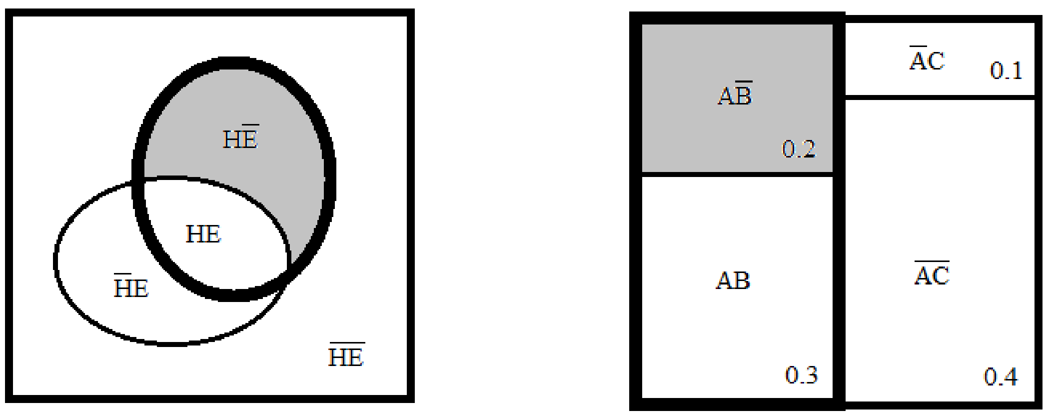

FIGURE 1: Muddy Venn Diagrams

Let us now consider how one's probability distribution could change in response to new facts coming to light. Learning new facts can be represented by moving mud around in the diagram. Here are two ways that one could move mud around, that each correspond to a particular way of ruling out the possibility that $A\bar{B}$. The first is that one could remove all the mud from the $A\bar{B}$ region and distribute it to the rest of the diagram while preserving facts about the *proportions* of the remaining mud (as in the left-hand side of Figure 2). The probabilities of AB, $\bar{A}\bar{C}$, and $\bar{A}C$ are 0.3, 0.4, and 0.1, respectively, and so to maintain the 3:4:1 ratio, we distribute probability so that p(AB) = 0.375, p($\bar{A}\bar{C}$) = 0.5, and p($\bar{A}C$) = 0.125. (This is equivalent to "removing" the $A\bar{B}$ mud and renormalizing, as van Fraassen describes the procedure.) This procedure preserves the probability ratio of the remaining possibilities, and corresponds to the familiar updating rule known as (classical) *conditionalization*. Indeed, Bayes' Rule is a formal characterization of the operation of this procedure when we learn some evidence E (i.e. rule out $H\bar{E}$ and $\bar{H}\bar{E}$). So, the first way to move mud around is to *maintain the ratios between the unconditional probabilities of the remaining options*. One obvious effect of conditionalizing by eliminating $A\bar{B}$ is to

[13] There may be good philosophical reasons to think of conditional probability as primitive (see Hájek 2003 and Pruss (2012)). Nonetheless, this won't matter to our discussion.

lower the probability of A relative to $\bar{A}$: since one of the "A" regions is eliminated, and the remaining regions retain their ratios with each other, the probability of A can only go down. This is what happened in the case of the urns: T12 was (nearly) ruled out, and the mud previously assigned to that region was redistributed to the remaining hypotheses in proportion to their previous probability. Similarly, to borrow another example from the skeptical theism literature, the same result should obtain when you find yourself in the following situation.[14] You are not sure whether your friend is in town (A) or out of town ($\bar{A}$), but if he is in town there are a limited number of possibilities. You check the concert hall ($A\bar{B}$), and he is not there. In this situation, the probability you assign to "he is in town" ought to decrease. You've learned that he is not at the concert hall, and conditionalized on that fact. *Purely incidentally*, you've also learned the (material) conditional <If he is in town, he is not at the concert hall>.

A second way that one could move mud around in the diagram while ruling out the possibility that $A\bar{B}$ is to take all of the mud from $A\bar{B}$ and move it to AB (as in the right-hand side of Figure 2). This procedure preserves the ratio between p(A) and p($\bar{A}$), since mud is moved around only within the "A" region, not between the two regions. But it does not preserve the ratio between AB and any of the other regions: the probability of AB increases relative to the probability of $\bar{A}\bar{C}$, for example. What type of learning, if any, could result in such a change? One example is learning the indicative conditional <If A then B> in a situation in which learning this conditional is irrelevant to the probability of its antecedent. For example, you assign equal probability to the hypothesis that your friend is in town (A) and the hypothesis that he is out of town ($\bar{A}$). There are five coffee shops in town, three Pete's and two Starbucks, and knowing nothing else, you assign equal probability to his being at each (with AB representing his being in town at a Pete's and $A\bar{B}$ his being in town at a Starbucks). You then learn that he hates Starbucks, so if he's in town, he won't be there – therefore, you can rule out $A\bar{B}$. Intuitively, though, learning this fact shouldn't make you think it more likely that he is out of town. Whereas the first procedure captured updating by conditionalization, this procedure captures *updating on an indicative conditional* without lowering the probability of the antecedent. In both cases, you've ruled out a version of one of the more general hypotheses: but in the first case, you've ruled it out in such a way as to make the general hypothesis less likely, and in the second, you haven't.

[14] Rowe (2004).

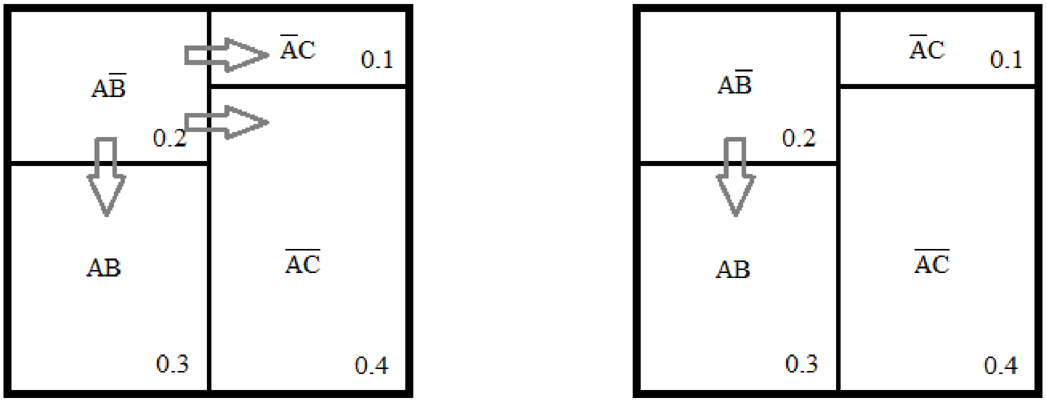

FIGURE 2: Two responses to ruling out $A\bar{B}$

Cell	Before	Change	After		Cell	Before	Change	After
$A\bar{B}$	.2	-.2	0		$A\bar{B}$	.2	-.2	0
AB	.3	.2(.3/.8) = .075	.375		AB	.3	.2	.5
$\bar{A}B$	.1	.2(.1/.8) = .025	.125		$\bar{A}B$	.1	0	.1
$\bar{A}\bar{B}$	.4	.2(.4/.8) = .1	.5		$\bar{A}\bar{B}$	.4	0	.4
SUM	1	0	1			1	0	1

FIGURE 3: Two ways of redistributing credences

How should this second kind of updating be formalized? And, moreover, how do we know when learning a conditional <If A then B> is a case of learning that $A\bar{B}$ is false that can be handled by conditionalizing on not-($A\bar{B}$), and when it is a case in which the probability of the antecedent ought to be preserved? That is, how do we know when we're in an updating situation in which we ought to set p($A\bar{B}$) = 0 while preserving the ratios of unconditional probabilities, and when we're in an updating situation in which we ought to set set p($A\bar{B}$) = 0 while preserving the antecedent probability of A? Unfortunately, there is currently no consensus.[15] Bradley (2005) argues that updating on a conditional without changing the antecedent probability can be modeled using what he calls "Adams conditioning," a special case of

[15] Douven/Romeijn (2011: 6) point out that it is not merely that there is no consensus, but that the question has received very little attention in the literature.

Jeffrey conditioning. He adds that how to update in a particular case requires a judgment call about "the epistemic standing of both our conditional and unconditional beliefs." Douven and Romeijn (2011) expand on Bradley's suggestion by suggesting a procedure for deciding which probability facts to preserve that takes into account which probability facts are more "epistemically entrenched." Lukits (forthcoming) argues for a rule known as MAXENT, which delivers the verdict that when updating on a conditional, the probability of the antecedent is preserved if the antecedent is casually independent of the consequent. And Bovens and Ferreira (2010) claim that standard Bayesian conditionalizing allows us to retain the prior probability of a conditional's antecedent as long as we include the fact that we learned the conditional in the set of facts we conditionalize on (this is a standard solution to the Monty Hall problem).

Despite the lack of consensus, three things are clear. First, there *are* cases of learning a conditional in which the probability of the antecedent ought to be preserved. Second, updating on a conditional <If A then B> in these cases cannot be modeled by conditionalizing exclusively on the material conditional (or on its truth-functional equivalent not-($A\bar{B}$)): either we will have to take into account additional information or we will have to use a different updating rule. Finally, knowing whether we are in one of these cases – whether the probability of the antecedent ought to be preserved – requires more than knowing which conditional we are updating on: it requires knowing which antecedent probabilities we are more committed to or how we came to receive the information that the conditional is true.

Is some scientific progress best modeled as learning a conditional while preserving the probability of its antecedent? One case of this might be that of a scientist adopting a research program and trying to discover some of its commitments, on the assumption that it is true and for reasons independent of what its rivals can explain. This describes Perrine and Wykstra's physicist "Grain," who gathers independent evidence (e2) for his conditional claim that if Granularism is true, then a particular version of it must be true. Learning which precisification of a general theory is the most plausible needn't always make the general theory less likely relative to the alternatives.

Now we can see what the key question for this debate is: when we take ourselves to learn that naïve theism is false – because we learn the data of good and evil and conclude that a theodicy won't work – should we conditionalize on not-($T\bar{S}$), thereby preserving the ratios between the remaining unconditional probabilities, or should we update on the conditional <If T then S> in such a way as to preserve the probability of T?[16] What those who reject theodicy (everyone in this debate) agree to is *at*

[16] Just to be clear: the *data* itself isn't "naïve theism is false." The data is some set of facts (D plus the supposed failure of theodicy) which has very low probability on naïve theism (I assume for simplicity: no probability), and higher probability on the remaining hypotheses, probability that is equal for all three hypotheses. This mirrors the structure of the above examples, in which the data were "I drew a yellow ball," "I checked the concert and my friend wasn't there," "a reliable third-party said my friend hates Starbucks," and e1; this data has low or no probability on

least that no or very low probability should be assigned to $T\overline{S}$ (let's assume no probability, for the sake of argument). But should that probability be re-distributed evenly across the remaining space, thus preserving the antecedent ratio between TS and N; or should it be re-distributed to T, thus preserving the antecedent ratio between T and N? The answer to this question depends on how exactly to characterize the data of good and evil, the failure of theodicy, and, perhaps, how we've come to learn these facts. Is what we take ourselves to have learned appropriately characterized as supporting not-($T\overline{S}$), or appropriately characterized as supporting the conditional <If T then S>?

Draper's urn example suggests that what we learn is exactly that a certain alternative (non-skeptical theism) is off the table – and if this is all we learn, then he is right that according to standard Bayesianism the evidence from evil supports theism over atheism, even when skeptical theism is on the table. Perrine and Wykstra's Granularism example suggests that what we learn is the conditional <If theism, then skeptical theism>. Notice that Perrine and Wykstra take Theo's absorption of the facts to primarily be a way of refining his theory. *Within* theism, he is considering what the best hypothesis is: he is considering, if theism is true, in what way is it most likely to be true? And he learns that skeptical theism is the best version of theism. If this is all *we* learn, then evil supports skeptical theism over other brands of theism but does not support atheism over theism.

We've already mentioned some situations that require conditionalizing on not-($A\overline{B}$) and some that require updating on the conditional <If A then B> while preserving the probability of A. While I cannot give a precise characterization of features that distinguish when ruling out $A\overline{B}$ ought to be characterized as conditionalizing on not-($A\overline{B}$) and when doing so ought to be characterized as updating on a conditional, here are a few general thoughts. The first is that, as several of the authors mentioned above point out, one ought to consider how one came to be in possession of the evidence that rules out $A\overline{B}$. What exactly do we learn *about our own learning* when we learn that naïve theism is false? In the scientific case, Grain started with a supposition and determined what followed from it: he asked, assuming Granularism is true, what's the best version of it? That the conditional was learned was independent of which general theory is true. This was also true in the second friend-locating case. In the urn case, one sampled randomly from the environment without regard to the assumptions of the general theories in question. This was also true in the first friend-locating case. The second thought (also pointed out by the above authors) is that when figuring out whether to conditionalize on not-($A\overline{B}$) or update on the conditional <If A then B>, one ought to think about which "probability facts" are more epistemically

---

the hypotheses H2, my friend is at the concert, my friend is at Starbucks, and a very coarse version of Granularism; and this data has higher probability on both T1 and T34, on both my friend is in town but not at the concert and my friend is out of town, both my friend is in town but not at Starbucks and my friend is out of town, and both fine-grained Granularism and Smoothism.

entrenched. For example, are we theists first and specific theists second, or does theism get its support from the initial plausibility of some of its specific versions?

I will close by pointing out that the general question discussed here reappears in a parallel debate: that about whether fine-tuning is evidence for theism. In particular, we might wonder whether the naturalist's invocation of the possibility of multiple universes blocks the claim that fine-tuning strongly supports theism over atheism. And we can approach this question by considering: when one learns that the physical constants are extremely unlikely to support life on single-universe atheism, is what one learns more accurately characterized as "If atheism, then multi-universe atheism" or as "Not single-universe atheism"? In both of these debates, the difference between the two ways of ruling out $A\bar{B}$ (naïve theism in the case of the argument from evil and single-universe atheism in the case of the fine-tuning argument) may explain why the proposed response (skeptical theism and multi-universe atheism, respectively) can seem ad hoc to those on the side of $\bar{A}$. The defender of A thinks of the data as supporting <If A then B>, whereas the defender of the claim that the datum undermines A thinks of the data as supporting not-($A\bar{B}$).

Examining the differences between Perrine/Wykstra and Draper reveals the way forward in the debate about whether skeptical theism undermines the argument from evil: we need to consider how to appropriately characterize the evidence, and which updating procedure to use to take it into account. Those who reject theodicy agree that we should assign no (or very low) probability to naïve theism in light of the evidence from evil and the history of theodicy. And if the evidence from evil and the history of theodicy is appropriately characterized as exactly "A theism that asserts we know what goods there are is incompatible with the level and type of evil in our world," (in short: naïve theism is incompatible with evil) then introducing skeptical theism as a potential hypothesis cannot blunt the force of the blow for the theist, since the initially most plausible specification of theism is ruled out without changing the relative probability of other versions of theism as compared with naturalism. However, if the evidence is appropriately characterized as "If God exists, then our knowing the goods there are is incompatible with the level and type of evil in our world," then hope remains for the skeptical theist.

Bibliography

Bovens, Luc and José Luis Ferreira (2010), "Monty Hall drives a wedge between Judy Benjamin and the Sleeping Beauty: a reply to Bovens." *Analysis* 70(3): 473-81.

Bradley, Richard (2005), "Radical Probabilism and Bayesian Conditioning." *Philosophy of Science* 72(2): 342-364.

Douven, Igor and Jan-Willem Romeijn (2011), "A New Resolution of the Judy Benjamin Problem." *Mind* 120(479): 637-70.

Draper, Paul, "Meet the New Skeptical Theism, Same as the Old Skeptical Theism." Forthcoming in Skeptical Theism: New Essays, eds. Trent Dougherty and Justin McBrayer (OUP).

van Fraassen, Bas C. (1989), Laws and Symmetry. Oxford: Clarendon Press.

Hájek, Alan (2003), "What Conditional Probability Could Not Be." *Synthese*, 137: 273-323.

Lukits, Stefan (forthcoming), "The principle of maximum entropy and a problem in probability kinematics." *Synthese*.

Perrine, Timothy and Stephen J. Wykstra, "Skeptical Theism, Abductive Atheology, and Theory Versioning." Forthcoming in Skeptical Theism: New Essays, eds. Trent Dougherty and Justin McBrayer (OUP).

Pruss, Alex (2012). "Conditional probabilities", *Analysis* 72: 488–491.

Rowe, William L. (2004), "Is Evil Evidence against Belief in God?" Chapter 1 of Contemporary Debates in Philosophy of Religion, eds. Michael L. Peterson and Raymond J. VanArragon. Pp. 3-13.