

The argument from so many arguments

Ted Poston*

May 6, 2015

There probably is a God. Many things are easier to explain if there is than if there isn't.

John Von Neumann

My goal in this paper is to offer a Bayesian model of strength of evidence in cases in which there are multiple items of independent evidence. I will use this Bayesian model to evaluate the strength of evidence for theism if, as Plantinga claims, there are two dozen or so arguments for theism.¹ Formal models are justified by their clarity, precision, and usefulness, even though they involve abstractions that do not perfectly fit the phenomena. Many of Plantinga's arguments are metaphysical arguments, involving premises which are necessarily true, if true at all. Applying a Bayesian account of strength of evidence in this case involves reformulating some of the arguments, but, even if a Bayesian shoe doesn't fit perfectly into a Leibnizian foot, Bayesian footwear is much more suitable to certain types of terrain, especially when the landscape requires encompassing the overall effect of multiple vistas. I believe that the Bayesian model I offer has significant utility in assessing strength of evidence in cases of multiple items of evidence. The model turns questions of the overall strength of multiple arguments into a simple summation problem and it provides a clear framework for raising more philosophical questions about the argument. I hope that this paper provides a model for many fruitful conversations about how to aggregate multiple items of evidence.²

*Thanks to Dan Bonevac, Robin Collins, Tim McGrew, Brad Monton, Alex Pruss, Bill Roche, and Richard Swinburne for helpful comments on an earlier draft. I am especially grateful to Branden Fitelson for multiple conversations that led to significant improvements in this paper.

¹(Plantinga 2007)

²My attitude to the Bayesian methodology is nicely captured by Paul Horwich's excellent article

1 A Bayesian model of multiple items of evidence

Bayesianism is well suited to evaluate the evidential impact of multiple items of evidence. It offers a clear account of evidential relevance and the evidential impact of evidence on theory. I begin by discussing several idealizations in order to offer a tractable Bayesian model of evidential impact. I then offer the results of the model and discuss some questions that arise from the model.

1.1 The Bayesian Framework

Bayesian models of the strength of evidence for a hypothesis are quite useful. In the following I discuss some aspects of the particular model I use. These aspects are idealizations but not every idealization involves a significant departure from reality. The model makes tractable a number of issues with respect to the overall confirmation of theism. Its fruitfulness justifies its use even if it doesn't faithfully capture all our intuitions about confirmation.

A Bayesian model requires that hypotheses and evidence can be given a prior probability that reflects the plausibility of the hypotheses and evidence prior to learning from experience. A probability is a real number inclusively between 0 and 1 that obeys the axioms of the probability calculus. I will take it for granted that we can make comparative claims about the probability of various competing hypotheses and the probability of some items of evidence on those hypotheses. This does not commit us to thinking that every proposition has a probability nor that the probability is some fixed real number.³ Moreover, I'll leave open the interpretation of a probability function. Personalists can read the forthcoming results as yielding at most consistency requirements for rationality. Rationalists can read the following results in a more robust sense. My task isn't to wade into these murky waters; rather I cast my net with an eye to putting on the table a clear model of the confirmation of multiple lines of argument.⁴

The specific probabilistic comparisons I will work with are comparisons of what is known as the *likelihoods*. The likelihood is the probability that some evidential claim is true given a specific hypothesis. If a fair coin is to be flipped then the probability that it lands heads is .5. The relevant likelihood is $\Pr(\text{heads} \mid \text{fair})$. The cases that

(Horwich 1993)

³Keynes, for instance, did not think that all probabilities were numerical. Rather on his view probabilities can be comparative. See (Keynes 1921)

⁴My argument does assume that the probabilities at issue are not physical propensities. I have in mind a notion of probability that is sensitive to a subject's evidence. For a recent defense of this notion of probability see (Maher 2006, 2010)

concern this paper are assigning comparative likelihoods of an evidential claim given theism and naturalism. Suppose our evidential claim is there is a physical universe whose natural laws allow for the development of embodied creatures. Call this claim “FT” for fine-tuning. The relevant likelihood ratio is the value of this:

$$\frac{Pr(FT|theism)}{Pr(FT|naturalism)}$$

The model I offer below is compatible with any estimate of these likelihoods. When we get to the model I will discuss the results on different assumptions about the values for the relevant likelihoods.

Another assumption I make concerns how many different hypotheses are at issue. I shall work with only two: the hypothesis of theism and the hypothesis of naturalism. Theism is the hypothesis that there is a omniscient, omnipotent, and omnibenevolent being. This is the hypothesis that there is a perfect being. I shall refer to the theistic hypothesis as ‘T’. Naturalism is the hypothesis that there is no such being or any similar beings, that the Democritus was right when he said that “it’s just atoms and the void.” For the purposes of the following Bayesian model, I assume the useful falsehood that theism and naturalism are mutually exclusive and exhaustive. Thus, naturalism will be represented as ‘¬T.’ Two reasons justify this useful falsehood. First, the confirmation of a hypothesis is a matter of how well it beats its conceived rival hypotheses. Given some evidence and a field of hypotheses, we ask ourselves how much does this evidence confirm one hypothesis over its competitors. The problems that arise with the catchall hypothesis are rightfully often ignored in a specific context in which a field of hypotheses are at issue. Second, my overall task is to model the evidential impact of multiple lines of argument. The model I offer is more tractable given the assumption that theism and naturalism are mutually exclusive and exhaustive options.

1.2 Evidential relevance & strength of evidence

According to the Bayesian account of evidential relevance, an item of evidence, e , is evidentially relevant to a hypothesis h just in case $Pr(h | e) \neq Pr(h)$. If e raises $Pr(h | e)$ then e is evidence for h and if e lowers $Pr(h | e)$ then e is evidence against h . Bayes’s theorem relates $Pr(h | e)$ with three other probabilities: (i) $Pr(e | h)$, called the ‘likelihood of the hypothesis’, (ii) $Pr(h)$, the prior probability of the hypothesis, and (iii) $Pr(e)$, the prior probability of the evidence. Bayes’s theorem states,

$$Pr(h | e) = \frac{Pr(e|h)Pr(h)}{Pr(e)}$$

In our case we are interested in comparing the confirmation of multiple items of independent evidence to theism and naturalism. Thus we can work with the odds form of Bayes’s theorem, which is as follows:

$$\frac{Pr(T|e)}{Pr(-T|e)} = \frac{Pr(T)}{Pr(-T)} \times \frac{Pr(e|T)}{Pr(e|-T)}$$

The odds form follows from Bayes’s theorem because both posterior probabilities are divided by the same quantity— $Pr(e)$. The usefulness of this form is that it allows us to directly compare their relative confirmation by examining the ratios of the priors and the likelihoods. The odds form isolates two factors in the judgement over the posterior probabilities: (i) prior belief and (ii) evidential strength. In the following I focus only on the relative strength of evidence, i.e., the ratio $\frac{Pr(e|T)}{Pr(e|-T)}$. This abstracts away from the influence of prior belief. I am interested in a measure of evidential strength that is not sensitive to differences in prior beliefs. Thus, measuring strength of evidence should only concern the relevant likelihood ratios.

1.3 Royall’s case

The statistician Richard Royall defends the claim that a likelihood ratio of 8 is ‘pretty strong’ evidence.⁵ Royall provides a natural case to ground the evidential strength of a likelihood ratio of 8. Suppose we have two urns. Urn 1 contains only white balls and urn 2 contains half white and half black balls. We select a ball from an urn, record its color, place it back in the urn, and thoroughly mix the contents of the urn. Suppose you begin to draw balls from the urns and you draw three successive white balls. It is natural that this is pretty strong evidence that the urn contains only white balls. The likelihood ratio in this case is $2^3 = 8$.

We can use Royall’s case to provide a natural anchor for any value to the likelihood ratio. Where b is the number of successive white balls then the likelihood ratio in favor of white over half white is $\frac{1}{.5^b} = 2^b$. Royall provides the following values for b and corresponding likelihood ratios.

Table 1: Number of successive white balls (b) corresponding to values of a likelihood ratio (LR)

LR	10	20	50	100	1000
b	3.3	4.3	5.6	6.6	10

⁵(Royall 1997, 12) Thanks to Branden Fitelson for pointing me to the relevance of Royall’s discussion

A nice feature of Table 1 is that we can go back and forth between a likelihood ratio and a sequence of white draws. Suppose we find that the evidence has a likelihood ratio of 20. What is the natural understanding of the strength of evidence corresponding to a likelihood ratio of 20? It's strong evidence, but how strong? By inspection the table shows us that a LR of 20 is the equivalent of selecting 4.3 white balls in sequence. If our evidence just consists in the fact that 4 white balls were selected in Royall's set up then that is significant evidence that we are selecting from the white only urn.

Royall's case is useful for associating a LR value with a more natural case. But the specific details of his case are not ideal for my purpose since we can get conclusive evidence that we are selecting from urn 2 if one black ball is selected. This feature of Royall's original case is easily avoided by changing the contents of the urn.⁶ Instead of urn 1 containing all white balls, let us suppose it contains $\frac{2}{3}$ white and $\frac{1}{3}$ black. Moreover, suppose urn 2 has $\frac{1}{3}$ white and $\frac{2}{3}$ black. Then $\Pr(\text{white} \mid \text{urn 1}) = .66$ and $\Pr(\text{white} \mid \text{urn 2}) = .33$. The relevant likelihood ratio in this case is 2. All the desirable features of Royall's case are maintained. If we observe 5 white balls in sequence then the likelihood ratio is $2^5 = 32$. This is very strong evidence that we are selecting from urn 1. In the following I will use this modified version of Royall's case.

Suppose that all you learn is that white_i and white_j , that is on the i^{th} and j^{th} draws white balls were selected. How strong of a case is that for urn 2? The relevant likelihood ratios are

$$\frac{\Pr(\text{white}_i \mid \text{urn1})}{\Pr(\text{white}_i \mid \text{urn2})} = 2$$

and

$$\frac{\Pr(\text{white}_j \mid \text{urn1})}{\Pr(\text{white}_j \mid \text{urn2})} = 2.$$

It is natural to think that white_i and white_j *independently confirm* that Urn 1 was selected. Since the selections from the urns are not affected by previous draws, we can take multiple white draws as each independent evidence for urn 1. In this case we calculate the cumulative power of multiple white selections according to the following:

$$\frac{\Pr(\text{White}_{(1,2,\dots,n)} \mid \text{Urn1})}{\Pr(\text{White}_{(1,2,\dots,n)} \mid \text{Urn2})} = \frac{1}{.5^n} = 2^n$$

This formula tells us that when we have n white draws the power of the evidence is 2^n . Because the likelihood ratio for black balls is the multiplicative inverse of the likelihood ratio for white balls, we can also easily accommodate the number, j , of

⁶Thanks to Tim McGrew for pointing this out.

black balls selected by using 2^{n-j} . Thus where our evidence is that 10 white balls were selected and 5 black balls were selected then the relevant likelihood ratio is 2^5 . That is strong evidence that we are selecting from Urn 1.

1.4 Other features of Royall's case

Royall's case is useful for introducing a number of probabilistic features pertaining to strength of evidence in cases of multiple items of evidence. Above we saw that because a specific draw of a ball from an urn is unaffected by the previous draws, we can easily combine the overall strength of evidence for all the draws. This feature is sometimes described under the heading of 'independent evidence.' But note that in Royall's case and the modified version the draws are not *unconditionally independent* from each other. The probability that white is selected on draw 2 *is* influenced by whether or not white was drawn on the first draw. The reason for this is that white on 1 provides evidence that urn 1 was selected and thus changes the probability that white is selected on draw 2.

Let us put these points using the formal model. Let ' W_i ' stand for the claim that 'a white ball is selected on the i^{th} draw' and ' U_j ' stand for the claim that 'Urn number j was selected.' The unconditional probability that white is selected on 1 is

$$Pr(W_1) = Pr(W_1 | U_1) \times Pr(U_1) + Pr(W_1 | U_2) \times Pr(U_2) = (0.66 \times 0.5) + (0.33 \times 0.5) = 0.5$$

Similarly, the unconditional probability that white is selected on draw 2 is 0.5. But the probability that white is selected on draw 2 given that white was selected on draw 1 is not 0.5. Rather it is the following:

$$Pr(W_2 | W_1) = Pr(U_1 | W_1)Pr(W_2 | W_1 \& U_1) + Pr(U_2 | W_1)Pr(W_2 | W_1 \& U_2)$$

We can use Bayes's theorem to determine ' $Pr(U_1 | W_1)$ ' and ' $Pr(U_2 | W_1)$ '. The values here are $\frac{2}{3}$ and $\frac{1}{3}$ respectively. Next we use the features of the modified version of Royall's case to determine the remaining values. $Pr(W_2 | W_1 \& U_1) = \frac{2}{3}$ and $Pr(W_2 | W_1 \& U_2) = \frac{1}{3}$. Thus, $Pr(W_2 | W_1) = \frac{5}{9}$. The formal model tracks our natural judgement that the selection of white on 1 is evidence that white will be selected on 2.

We noted above that in Royall's case we can combine the selections of white balls easily to get the overall evidential strength of n white balls. What probabilistic feature of Royall's case is responsible for the natural judgment that W_i and W_j are

confirmationally independent regarding the hypothesis that urn 1 was selected. The key feature is that the selection of an urn *screens off* the results of previous draws. The screening off condition is the following, where $j > i$

1. $Pr(w_j | Urn1 \& w_i) = Pr(w_j | Urn1)$
2. $Pr(w_j | Urn2 \& w_i) = Pr(w_j | Urn2)$

Claims (1) and (2) hold in Royall's model because once the urn is selected, the results of previous draws don't change the probability of any subsequent draw.

The key features of Royall's case are the following. First, given a likelihood ratio we can find a corresponding number of white balls. This grounds judgements about the evidential strength of various likelihood values. Second, the individual selections are not *unconditionally independent*. Third, the individual selections are probabilistically independent regarding U_1 and U_2 . The importance of probabilistic independence regarding H is that it ensures that we can combine (in a way to be discussed shortly) multiple items of evidence. It is important to realize that the crucial notion of independence is not unconditional independence but independence regarding H. When it comes to modeling multiple items of evidence, each item of positive evidence for a hypothesis may increase the probability that the hypothesis is true and thus lead us to expect more positive evidence. Even so, if the evidence is independent regarding H then it does not affect a natural way of combining multiple items of evidence. For reasons that I will not go into the second and third feature of Royall's case uniquely picks up the log likelihood ratio.⁷

1.5 The log-likelihood ratio & independent evidence

Independent evidence for a theory is important because the confirmatory significance of each item of independent evidence is undiluted when added together together in a larger case for the theory. C.S. Peirce remarks on this feature of independent evidence. He writes,

Two arguments which are entirely independent, neither weakening nor strengthening the other, ought, when they concur, to produce a[n intensity of] belief equal to the sum of the intensities of belief which either would produce separately. (Peirce 1878; my brackets)⁸

Branden Fitelson provides a nice Bayesian account of the conformational significance of independent evidence.⁹ Two items of evidence are confirmationally inde-

⁷See (Fitelson 2001)

⁸(Fitelson 2001, 125)

⁹(Fitelson 2001)

pendent regarding H if and only if the support each provides for H is independent of whether the other piece of evidence is already known. That is,

Definition of confirmational independence: E_1 and E_2 are (mutually) confirmationally independent regarding H according to c if and only if both $c(H, E_1|E_2)=c(H, E_1)$ and $c(H, E_2|E_1)=c(H, E_2)$.¹⁰

$c(H, E_1)$ is the degree of confirmation that E_1 provides for H. $c(H, E_1|E_2)$ is the degree of confirmation that E_1 provides for H given that E_2 is known. This definition captures the idea that the confirmational significance of independent evidence is unaffected by whether the other evidence is known. Fitelson shows that the log likelihood measure (or an ordinally equivalent measure) satisfies natural principles about evidential independence such as Pierce's claim that when two arguments are entirely independent then when they both occur they should produce an level of confidence equal to the sum of the levels of confidence each would produce separately.¹¹ Fitelson interprets this Peircean claim as the following:

(A) If E_1 and E_2 are confirmationally independent regarding H according to c , then $c(H, E_1 \& E_2)=c(H, E_1)+c(H, E_2)$.

This claim allows us to add together the confirmational significance of items of independent evidence. Thus, where E_1 and E_2 are confirmationally independent regarding H according to c , we can determine the total confirmation E_1 and E_2 offer to H by the following:

$$\log\left(\frac{E_1|H}{E_1|\neg H}\right) + \log\left(\frac{E_2|H}{E_2|\neg H}\right)$$

An attractive feature of using logarithms is that it turns multiplication into addition. Thus, where the evidence meets the screening off condition we can simply add together the evidential effect of multiple items of evidence.

2 Models for multiple items of evidence

In the following I provide two models for the evidential strength of multiple items of evidence. The models vary by the relevant likelihood ratio. In the first model I assume that the likelihood ratio for each item of evidence favors theism by a factor of

¹⁰(Fitelson 2001, 125)

¹¹see (Fitelson 2001, 125)

10. Using base-10 log, this gives us a log likelihood ratio of 1.¹² In the second model I assume that the likelihood ratio for each item favors theism by 10%. I also discuss ways in which the models can be used to account for evidence against theism.

2.1 Strong independent evidence

The initial model is one in which the positive evidence for theism is an order of magnitude greater on theism than on naturalism. For each e_i , $\log\left[\frac{\text{Pr}(e_i|T)}{\text{Pr}(e_i|\neg T)}\right] = 1$ In Royall's model this is strong evidence; it is equivalent to the selection of 3.3 white balls.

The assumption that the log likelihood ratio for each item of evidence for theism is not completely indefensible. Richard Swinburne thinks that if we can determine an agent's beliefs, desires, and goals then we can determine with significant probability that an agent will do some thing.¹³ In the case of theism Swinburne holds that we can make reasonable determinations of what kinds of action a perfect being will bring about. A morally perfect being will bring about the best action if there is one and it is within the scope of his knowledge and power. An omniscient and omnipotent morally perfect being will therefore bring about the best action if there is one and will not perform a bad action. To the extent there are many different incompatible good actions with no best, we can reasonably determine that a perfect being will bring about one out of those n good options. Thus, Swinburne thinks we can assign a probability to the claim that God will bring about some good option such as the creation of other beings of limited knowledge, power, and goodness. Since we are interested in purely the comparative probabilities, it is not unreasonable to think that there is an order of magnitude difference with respect to whether there would be finite beings of intentional power and goodness given theism as opposed to given naturalism. Whether this holds for multiple items of independent evidence is an open question, but I shall assume so in this model.

A nice feature of this model is that it is a simple head count to determine the overall strength of evidence for theism. If Plantinga is right that there are 24 arguments for theism then, on this model, the overall evidential strength of theism is $c(T, e_1, e_2, \dots, e_{24}) = 24$. Given this value, we can then go back to Royall's case to determine the equivalent number of white balls.

We need to do a bit of work first to find that number. First, consider table 2 that looks at the strength of evidence provided by the first 15 white balls. We see that

¹²Natural logs are more prevalent in mathematics but base-10 logs are easier to work with. Nothing hangs on the choice of a base.

¹³(Swinburne 2004, 112-123)

15 white balls in sequence only gives us a LLR of 4.5. In order to find the equivalent number of white balls for a LLR of 24 we have to continue the count. Because the numbers get large fast, I have shortened the results in table 3.

Table 2: Comparison of Royall’s likelihood ratios with the log likelihood ratio

# of successive white balls	Likelihood ratio (LR) 2^n	log likelihood ratio (LLR)
1	2	0.301029995663981
2	4	0.602059991327962
3	8	0.903089986991944
4	16	1.20411998265592
5	32	1.50514997831991
6	64	1.80617997398389
7	128	2.10720996964787
8	256	2.40823996531185
9	512	2.70926996097583
10	1024	3.01029995663981
11	2048	3.31132995230379
12	4096	3.61235994796777
13	8192	3.91338994363176
14	16384	4.21441993929574
15	32768	4.51544993495972

By inspection on table 3 we can see that the evidential strength of 24 independent arguments of theism each with a LLR of 1 is the equivalent of approximately 80 white balls. Using the modified version of Royall’s case, even if we have some black balls in the sequence, the difference of white to black is 80. A LLR of 24, then tells us that the evidential case for theism is as strong as the case in which we have 80 more white balls than black. That is overpowering evidence that urn 1 is selected.

This table shows up how to translate a log likelihood ratio back into Royall’s case. We can see by inspection that if there are 24 independent arguments for theism then that is like a string of 80 successive white balls. That is overwhelming evidence in favor of an all white urn.

In general, if there are n independent arguments for theism (where each is favorable to theism by a log likelihood ratio of 1) then we can determine the corresponding number of successive white balls according to the following equation: $b = \frac{n}{0.30102999566398}$. Table 4 produces the value for b for various determinations of how many independent arguments for theism there are.

Table 3: Comparison of successive white balls with log likelihood ratios

# of successive white balls	log likelihood ratio
1	0.301029995663981
2	0.602059991327962
3	0.903089986991944
4	1.20411998265592
5	1.50514997831991
10	3.01029995663981
20	6.02059991327962
30	9.03089986991944
40	12.0411998265592
50	15.0514997831991
60	18.0617997398389
70	21.0720996964787
80	24.0823996531185

Table 4: Comparison of # of independent arguments for theism with # of successive white balls

# of independent arguments for theism	# of successive white balls
5	16.6096404744368
10	33.2192809488736
15	49.8289214233104
20	66.4385618977472
25	83.0482023721841

A fruitful aspect of using this model is that it is easy to determine the evidential strength of multiple arguments. Suppose there are 5 independent arguments then the evidential case is the equivalent of the selection of 16.6 white balls. The evidential strength of multiple arguments is directly proportional to the number of independent lines of evidence.

Moreover, we can use the odds form of Bayes's theorem to determine the relevant difference in prior belief that would be required for a person unmoved by the evidential case. To be unmoved by the strength of evidence requires that $\frac{Pr(T)}{Pr(-T)}$ must be near the reciprocal of the likelihood ratio. If someone is unmoved by strength of evidence having an LLR of 24 then the difference among the priors must be $\frac{1}{10^{24}}$.

This model can also accommodate the evidence against theism. Assuming that

each item of negatively relevant evidence has a LLR of 1 in favor of atheism, we simply subtract the number of items of evidence against theism from the number of items of evidence for theism. The final count gives us the new LLR. If, for instance, there are 5 arguments for theism and 2 arguments against, then, assuming the LLRs are same, the evidential case as a LLR of 3. Consulting table 3, we see that a LLR of 3 is the equivalent of the selection of 10 white balls. That is significant evidence.

The model can also accommodate items of evidence that have different LLRs. Many philosophers hold that the primary item of evidence for theism is facts about the distribution and quality of evils we observe in our world. It is not unreasonable to think of this as very powerful evidence against theism. One can easily express the power of this evidence in terms of some multiple of a LLR of 1. If we assign the problem of evil a LLR of 4 then this expresses the judgement that the strength of evidence that evil has is like the selection of 13 white balls. The overall tally procedure is the same.

2.2 Weak independent evidence

Let us examine a model in which the relevant LLRs are only slightly in favor of theism. For each item of evidence let us suppose that the theist as a slightly more plausible story than the naturalism. How shall we interpret probabilistically the idea that a hypothesis as a slightly more plausible account of the evidence than its competitor? Thinking purely about the likelihood ratio, a favorable ratio of a 10% captures the idea that the evidence only weakly favors H over \neg H.

In the following I work with this: $\frac{Pr(e_i|T)}{Pr(e_i|\neg T)} = \frac{a}{\frac{9}{10}a} = \frac{10}{9}$. A likelihood ratio that favors T over \neg T by a 10% gives us a LLR of 0.046.¹⁴ At this value it would take approximately 22 independent lines of evidence to equal the confirmatory power of 1 white ball.

What we learn from this is that the values of the LLRs are crucial. We can easily accommodate differing LLRs by thinking about how powerful the evidence is in terms of Royall's model. Does the individual evidence have the power of a selection of 1 ball, 2 balls, 3 balls, and so on? In each case, we find the corresponding LLR and add together the cumulative effect of the evidence. As noted above, we can add together different values of the LLRs. If one piece of evidence is weak but another piece is strong then we simply add the LLRs and find the corresponding number of white balls.

Another fruitful feature of these models is that we can estimate a person's LLRs

¹⁴Using log base 10.

for the evidence if we know the ratio of their prior beliefs and the ratio of their posterior beliefs. If a person thinks that theism and naturalism have roughly the same prior and yet they remain agnostic given the evidence, then the overall impact of all the evidence has an LLR of 0. A different person who thinks that the evidence is every so slightly in favor of theism and whose difference in prior belief is a wash, may well be in a position to engage in Pascalian reasoning. I find these Bayesianism models fruitful if only for clarifying a person's judgements about the relevant likelihood ratios. In the following I shall examine some other fruitful aspects of these models.

3 Independence & Evidence

In this final section I tackle two questions. First, the model that I've used requires that multiple items of evidence meet the screening off condition. This is the following condition:

$$(S) \frac{Pr(e_i|T)}{Pr(e_i|\neg T)} = \frac{Pr(e_i|T \& e_j)}{Pr(e_i|\neg T \& e_j)}, \text{ for each } e_i, e_j.$$

(S) requires that once we specify the hypotheses and the previous items of evidence the probability for each particular item of evidence is not affected by this specification. As noted in section 1 (S) expresses the idea that the evidence is independent regarding h , and this is different from the idea that the evidence is unconditionally independent. It is crucial for a large evidential case that the evidence is not 'double counted'. This can happen if one piece of evidence entails the other. For instance, if one argument for theism proceeds from the claim that there is intentionality and another argument proceeds from the claim that there are moral beliefs then these arguments cannot both be added together because they violate condition (S). Since moral beliefs implies intentionality the probability of the latter given the former is 1. It is a fruitful project to go through Plantinga's two dozen or so arguments for theism and figure out which arguments violate condition (S). Trent Dougherty and I argued in a previous paper that multiple design arguments violate condition (S) and so the natural theologian cannot offer both in a overall cumulative case argument.¹⁵

Condition (S) forces the natural theologian to think about independent families of theistic arguments. It is an open question how many of these arguments are independent. It strikes me that the major theistic arguments have an interpretation in which they are conditionally independent. The ontological, cosmological, teleological, and moral arguments plausibly meet condition (S). The central datum in the ontological

¹⁵(Dougherty and Poston 2008)

argument that *it appears to be possible that there is a perfect being* does not entail the central datum in any of the other arguments. Those are respectively: *it appears as if the universe began to exist, that the conditions required for the development of embodied, rational creatures appear to be fine-tuned, that there are embodied, rational creatures who are sensitive to moral aspects of reality.*

When we explicitly think about the classic theistic arguments we meet a difficulty in how these arguments should be captured by a Bayesian model. The transcription I just gave of these classic four arguments for theism shifted the statement of the data from principles that are necessarily true, if true at all, to principles about appearances or seemings. A crucial problem with the Bayesian model I've used is that it requires that the data be given a non-zero probability on naturalism. This does require some reformulation. The crucial question is whether it is worth the cost? The answer to that question depends on future research. First, it needs to be determined how faithful the Bayesian reformulation is to the original metaphysical arguments. Second, it needs to be determined if there are other models for determining strength of evidence in cases in which the data itself either has a probability of 1 or 0. The value of the Bayesian model I've presented lays in clearly putting on the table these questions and giving us a clear model of arguments that do meet the relevant constraints.

4 Conclusion

I've argued that a fruitful approach to Plantinga's suggestion that there are two dozen or so arguments for theism is to work with a Bayesian model of the significance of multiple arguments for theism. We've seen that on some defensible assumptions about the relevant likelihood ratios, multiple arguments for theism provide significance evidence for theism. We've also seen that the values of the likelihood ratios are crucial. This requires that one defend not only the claim that the evidence is predicted by theism but that it is also not so strongly predicted by naturalism. Finally, I've looked at the crucial screening off condition and argued that the evidence needs to meet this condition to be summed. We've also looked at a consequence of this model that the relevant evidence must have a non-zero probability on naturalism.

References

- Dougherty, T. and Poston, T. (2008). A user's guide to teleological arguments, *Religious Studies* 44: 99–110.

- Fitelson, B. (2001). A bayesian account of independent evidence with applications, *Philosophy of Science* **68**(3): 123–140.
- Horwich, P. (1993). Wittgensteinian bayesianism, *Midwest Studies in Philosophy* **18**(1): 62–75.
- Keynes, J. M. (1921). *A Treatise on Probability*, Macmillan.
- Maher, P. (2006). The concept of inductive probability, *Erkenntnis* **65**: 185–206.
- Maher, P. (2010). Bayesian probability, *Synthese* **172**: 119–127.
- Plantinga, A. (2007). Two dozen (or so) theistic arguments, in B. Deane-Peter (ed.), *Alvin Plantinga*, Cambridge University Press, pp. 203–227.
- Royall, R. (1997). *Statistical Evidence: A likelihood paradigm*, Chapman and Hall.
- Swinburne, R. (2004). *The Existence of God*, 2nd edn, Oxford University Press.