

Selection Biases in Likelihood Arguments

Matthew Kotzen

ABSTRACT

Most philosophers accept some version of the requirement of total evidence (RTE), which tells us to always update on our complete evidence, which often includes ‘background information’ about how that evidence was collected. But different philosophers disagree about how to implement that requirement. In this article, I argue against one natural picture of how to implement the RTE in likelihood arguments, and I argue in favor of a different picture. I also apply my picture to the controversy over the so-called ‘Objection from Anthropic Bias’ to the fine-tuning argument, and argue that the Objection from Anthropic Bias fails.

- 1 *Introduction*
 - 2 *The Likelihood Principle*
 - 3 *The FTA*
 - 4 *Eddington’s Fish and the Requirement of Total Evidence*
 - 5 *LP* and the Objection from Anthropic Bias*
 - 6 *Problems with LP**
 - 6.1 *The line-drawing problem*
 - 6.2 *Firing squad cases*
 - 6.3 *Probabilistic dependence between I and the hypotheses*
 - 7 *My Proposal*
 - 8 *Other Views*
-

1 Introduction

It is a truism of scientific and philosophical methodology that, when we collect evidence, it can matter how that evidence was collected, and that ignoring ‘background information’ about how our evidence was collected can lead to irrational conclusions. For instance, when we interview people and ask them for their opinions about who will win the upcoming election, we are likely to

get different answers in predominantly Conservative areas than we are to get in predominantly Liberal areas; as a result, the conclusions that we reach on the basis of our interviews must take appropriate account of where and how those interviews were conducted.

However, what it means to ‘take appropriate account’ of how our data were collected is not always a straightforward matter. In this article, I’ll argue that one very natural and influential approach to taking background information into account is mistaken, and that it should be replaced with a rather different one.

The approach that I’ll defend is quite general, but it does have consequences for some particular controversies in epistemology and the philosophy of science. The controversy that I’ll be focusing most on is the controversy over the ‘Fine-Tuning Argument’ (FTA) and the ‘Objection from Anthropic Bias’ that some philosophers have raised against the FTA. As we’ll see, the Objection from Anthropic Bias takes the form of a charge that the FTA fails to take appropriate account of relevant background information; thus, different approaches to background information seem to be underlying and motivating each side of the debate. My approach entails that the Objection from Anthropic Bias fails; thus, if there is a problem with the FTA, the Objection from Anthropic Bias has not successfully identified it. But again, my argument has consequences beyond the Objection from Anthropic Bias and the FTA; I will be arguing for a general theory of how background information ought to be taken into account in a wide class of scientific and philosophical arguments.

2 The Likelihood Principle

All of the arguments that I’ll be considering in this article are ‘likelihood arguments’, which depend on a principle like

Likelihood Principle (LP): E is evidence for $H1$ over $H2$ iff $P(E|H1) > P(E|H2)$

In this article, I’ll be using conditional probabilities such as $P(E|H1)$ and $P(E|H2)$ to refer to the objective likelihood that some hypothesis (tenselessly) confers on some evidence. So, for example, $P(E|H1)$ will refer to the objective probability that E would occur given that $H1$ is true. Even if we (now) know that E occurred, that doesn’t entail that $P(E|H1) = 1$ as I’m understanding conditional probabilities in this article, since the truth of $H1$ may not have *objectively guaranteed* the truth of E . Since it is often legitimate to inquire about the evidential impact of ‘old evidence’,¹ we must not interpret $P(E|H1)$ as our actual, current, subjective credence in E , supposing $H1$ to be true.

¹ For instance, suppose that I know that Mary won the lottery three times in a row (E), and I’m just now considering what that fact provides me with evidence for. Suppose $H1$ is that the lottery

The idea behind LP, then, is that if one hypothesis makes E objectively more likely than another hypothesis, then the fact that E actually does occur is some evidence for the first hypothesis over the second. While there are certainly some philosophers who have raised doubts about the core idea behind LP,² that core idea has been extremely influential and is accepted in some form by nearly all so-called ‘Likelihoodists’ and ‘Bayesians’. I will take this approach to confirmation theory for granted throughout this article; thus, while I will entertain (and argue for) alternatives to LP itself, I will assume that some version of the ‘core idea’ behind likelihood arguments is correct. Thus, one can read my argument as endeavoring to establish that LP is not the best version of that core idea, and that a different version of the likelihood principle (one which I call LP**) is superior.³

3 The FTA

The FTA is a likelihood argument, which appeals to the intuition that a life-sustaining universe like ours was objectively much likelier to exist on the supposition that an intelligent designer created it than it was on the supposition that it was created by mindless natural processes. In the FTA, we use the following abbreviations:

E = The constants of the universe have life-sustaining values.

D = An intelligent designer set the values of the constants.

C = The values of the constants were set by mindless natural processes.

According to the FTA, $P(E|D) > P(E|C)$, so E is evidence for D over C by LP.

Why, according to the FTA, is $P(E|D) > P(E|C)$? First, it seems that $P(E|C)$ has to be very very low, since it seems to be extraordinarily unlikely that the values of the constants would just happen to fall in the narrow life-sustaining range if those constants were set by mindless natural processes; after all, the vast ‘majority’ of the metaphysically possible values that the constants could have taken are not life-sustaining.⁴ In contrast, if there is a designer who set

is fair and $H2$ is that the lottery is rigged in Mary’s favor. In considering the question of what evidential impact E has on $H1$ and $H2$, the relevant understanding of conditional probability is one where $P(E|H1)$ is low and $P(E|H2)$ is high (so that E is evidence for $H2$ over $H1$), even though I already know that E is true, and hence my subjective credence in E (conditional on anything) is 1. The most natural such understanding of conditional probability is the objective one described above, where $P(E|H1)$ refers to how objectively likely Mary was to win the lottery three times in a row if the lottery was fair. Some philosophers, however, do not think that these issues about old evidence are so easily sidestepped—see, e.g. (Monton [2006]).

² See, e.g. (Fitelson [unpublished]).

³ Thanks to Teddy Seidenfeld for pushing me on this point at the 2009 Formal Epistemology Workshop.

⁴ Of course I don’t really mean ‘majority’ here, since there are infinitely many metaphysically possible worlds in which the constants are set by mindless natural processes and are

the values of the constants, it seems to be rather more probable that that designer would intentionally see to it that the constants were life-sustaining, since universes containing life seem to be the sort of thing that designers would (or, at least, plausibly might) intentionally aim at.

In order for the FTA to go through, it's not required that $P(E|D)$ be high. In other words, it's not required that we think that an intelligent designer would be especially likely to set the constants to life-sustaining values. All that is required for the FTA to go through is that the constants are *more* likely to take on life-sustaining values if D is true than if C is true, that is, that a designer is *more* likely to set the constants to life-sustaining values than some mindless natural process is to result in constants with life-sustaining values. Since this has seemed very plausible, it has seemed plausible that $P(E|D) > P(E|C)$, and so, by LP, that E is evidence for D over C .

4 Eddington's Fish and the Requirement of Total Evidence

A case that is commonly invoked in discussions of LP and background information is one from Eddington ([1939]). Suppose that there are a bunch of fish in a pond, and I'm wondering whether the pond consists of mostly large fish (L) or mostly small fish (S). Suppose that I have a net that can catch only one fish at a time, and suppose that there are enough fish in the pond so that I am certain to catch one fish each time I throw my net into the pond.

Now, I throw my net into the pond at t and catch a large fish (E). According to LP, this is some evidence for L over S , since $P(E|L) > P(E|S)$, and that seems like the right result. But then suppose that I look down at my net and notice that the holes in my net are too large to catch small fish (I). Once I notice this, it would be epistemically irresponsible for me to ignore it and take myself to still have evidence for L since $P(E|L) > P(E|S)$. After all, to ignore this information would seem to be a violation of the requirement of total evidence (RTE), according to which we should always take account of the *strongest* statement of our evidence that we have available to us. And the RTE is well-motivated by lots of intuitive cases. For example, suppose that I know that a corn snake is slithering toward me. That is no evidence that I'm about to die, since I know that corn snakes don't ever hurt people. But if instead of

life-sustaining, and also infinitely many possible worlds in which the constants are set by mindless natural processes and aren't life-sustaining. So what I really mean is that there is some suitable measure over the space of possible worlds, and that the set of possible worlds in which the constants are set by mindless natural processes and are life-sustaining gets a much lower measure than the set of possible worlds in which the constants are set by mindless natural processes and are non-life-sustaining. This claim is analogous to the claim that in the vast 'majority' of possible worlds, the number of grains of sand in the universe is not positive and evenly divisible by 1,000,000,000.

considering the fact that there's a *corn* snake slithering toward me, I had instead considered only the (true) fact that there's a snake slithering toward me, that would be evidence that I'm about to die, since snakes harm people all the time. So, in order for me to come to the justified conclusion that I'm not about to die (rather than the unjustified conclusion that I am about to die), I need to take my evidence to be the stronger fact that there's a corn snake slithering toward me, rather than the weaker fact that there's a snake slithering toward me.

In Eddington's fish case, how should we respect the RTE and take the information about the size of the holes in my net into account? Clearly, we can't just compare $P(E|L)$ and $P(E|S)$ as LP recommends, since that comparison ignores I entirely. Rather, the relevant question has seemed to many to be whether $P(E|L \wedge I) > P(E|S \wedge I)$.⁵ In other words, the question is whether I am likelier to catch a large fish if the pond consisted of mostly large fish *and the holes in my net are too large to catch small fish* or if the pond consisted of mostly small fish *and the holes in my net are too large to catch small fish*.

Thus, the lesson of Eddington's case is often implicitly taken to be that LP needs to be amended (or perhaps just precisified) to something like:

Likelihood Principle* (LP*): Where I is our information about the process of observation, E is evidence for $H1$ over $H2$ iff $P(E|H1 \wedge I) > P(E|H2 \wedge I)$

In Eddington's case, given that I'm using a net with holes too large to catch a small fish, I'm just as likely to catch a large fish in my net regardless of whether L or S is true, so $P(E|L \wedge I) = P(E|S \wedge I)$. So, by applying LP* in Eddington's case, we get the correct result that E is no evidence for L over S .

5 LP* and the Objection from Anthropic Bias

Now, we are in a position to articulate the Objection from Anthropic Bias to the FTA. According to that objection, the FTA ignores a certain type of observation bias analogous to the bias present in Eddington's fish case, and (just as in Eddington's case) we violate the RTE by failing to take account of that bias. According to the Objection from Anthropic Bias, the relevant bias is constituted by the fact that intelligent observers like us can exist only in life-sustaining universes. But, the objection continues, it's no evidence for anything when we observe a precondition for our own existence, since that was *guaranteed* to happen. Thus, since the universe's being life-sustaining is a precondition for the existence of observers like us, we were guaranteed to

⁵ See, e.g. (Sober [2009]) and (Weisberg [2005]).

observe that the universe is life-sustaining, and hence the fact that the universe *is* life-sustaining can't be any evidence for D over C .

The application of LP* to the FTA is as follows: Just as I included the fact that the holes in my net are too large to catch small fish (on pain of violating the RTE) in Eddington's case, it seems that in the FTA, I should (on pain of violating the RTE) include the fact that I'm around at t (some time after my birth) to make observations, since that is part of the information that we have about the relevant process of observation. So the relevant question to ask in the FTA isn't whether $P(E|D) > P(E|C)$ (as LP recommended), but rather whether $P(E|D \wedge I) > P(E|C \wedge I)$ (as LP* recommends). But it's fairly obvious that if I includes the information that I'm around at t to make observations, then $P(E|D \wedge I) = P(E|C \wedge I) = 1$, since if I'm around at t to make observations, I'm certain to observe life-sustaining constants regardless of whether D or C is true. Thus, according to the Objection from Anthropic Bias, we can see that the FTA fails once we understand that respecting the RTE requires us to take into consideration our full information about the relevant process of observation, and once we accordingly reject LP in favor of LP*.

6 Problems with LP*

6.1 The line-drawing problem

According to LP*, when calculating likelihoods, we have to assume information I about the relevant process of observation 'on the right', i.e. by conjoining I to the expression on the right-hand side of the conditionalization bar. Moreover, in order to respect the RTE, I has to be our *total* evidence about the relevant process of observation; if we leave any information about the process of observation out, we are in danger of making the same mistake as I would if I took a corn snake slithering toward me to be evidence that I am about to die.

But this naturally leads us to wonder: How total should our 'total evidence' about the process of observation be? We learned from Eddington's fish case that I needs to include facts about the size of the holes in the net that I use. But another fact about the process of observation in Eddington's fish case is that the process of observation results in my catching a large fish. *That's* a fact about the process of observation too, so if the information that the net has large holes is included in I , it's not at all obvious why the information that the process of observation resulted in my catching a large fish shouldn't be included as well.

But this would be disastrous, and would completely trivialize LP*. For example, consider a 'bias-free' variant of the Eddington fish case where the net that I use has holes small enough to catch small fish, so that E really is evidence for L over S . In this case, if we include only the information that the net has small holes in I , then LP* gives us the right result that since

$P(E|L \wedge I) > P(E|S \wedge I)$, E is evidence for L over S . But if we also include in I the information that the process of observation resulted in my catching a large fish, then LP* gives us the *wrong* result that E is no evidence for L over S , since then $P(E|L \wedge I) = P(E|S \wedge I) = 1$; after all, given that my net has small holes *and that I catch a large fish*, I'm certain to catch a large fish regardless of whether L or S is true. More generally, *whenever* we include information in I which entails E , it will follow that $P(E|H1 \wedge I) = P(E|H2 \wedge I)$, so that E will never be any evidence for any hypothesis over any other hypothesis, which is absurd.

So, LP* leads us into a puzzle: How much information about the process of observation ought to be included in I ? We have to include in I the information that the net has large holes, but we can't include in I the information that the process of observation results in the collection of E . It's somewhat hard to see how to draw this line between the information about the process of observation that should be included in I and the information that shouldn't be included in I in a principled way, and the burden seems to be on the defender of LP* to do this. Call this the line-drawing problem.

There have been attempts to solve the line-drawing problem (though not presented in these terms) in the literature on fine-tuning. In particular, Sober ([2004], [2009]) and Weisberg ([2005]) have defended views about how much information ought to be included in I . Sober thinks that all of the information about the process of observation should be included,⁶ whereas Weisberg thinks that only *conditionals* about what you could not have observed should be included in I . I don't intend to enter into this debate here, as I think the debate is misguided. But the line-drawing problem is not a mere hypothetical one; it is a problem that philosophers who assume LP* have already spent time trying to solve.

6.2 Firing squad cases

A second problem with LP* is that there are cases which *seem* to be relevantly like the FTA, but where including information analogous to 'I'm around at t making observations' in I gets us the wrong results. This puts pressure on the defender of the Objection from Anthropic Bias to explain why 'I'm around at t making observations' *should* be included in I in the FTA, whereas the analogous information shouldn't be included in I in these cases.

Here's one such case:⁷ Suppose that I'm in front of a firing squad, and I know that the shooters are either going to intentionally miss me (D) or that

⁶ Sober doesn't specifically address the issue of whether the information that the process of observation resulted in the collection of E should be included in I . I assume that he doesn't think that it should, but I don't know how he solves the line-drawing problem.

⁷ This type of case was first introduced in (Leslie [1989]), and is discussed extensively in (Weisberg [2005]).

they are going to fire at random in my general direction (C). Intuitively, when I observe at t (a moment after the shooting) that I'm alive, that is some evidence for D over C , since I was much more likely to survive the firing squad on the assumption that the shooters intentionally missed me than I was on the assumption that the shooters shot at random (since it's fairly likely that at least one randomly shot bullet would fatally wound me). In other words, it's intuitive that my observation that I'm alive at t (E) is evidence for D over C precisely because $P(E|D) > P(E|C)$. But if LP* is true (and LP is false), then that probabilistic inequality is irrelevant. Rather, what is relevant is whether $P(E|D \wedge I) > P(E|C \wedge I)$. And if, as the Objection from Anthropic Bias holds, 'I'm around at t to make observations' should be included in I in the FTA, then it seems that that information should be included here too. But then $P(E|D \wedge I) = P(E|C \wedge I) = 1$, since I'm certain to observe that I'm alive at t given that I'm around at t to make observations, regardless of whether D or C is true. Thus, LP* leads to the implausible result that my observing at t that I'm alive is no evidence for D over C in the firing squad case.

Moreover, I don't think that this case essentially involves a situation where the relevant hypotheses affect my chances of surviving some ordeal, rather than my chances of coming to exist in the first place. Suppose that I know the following: Just after the egg that was to become me was fertilized, doctors took that fertilized egg and placed it in a petri dish. It's not known what happened after that—either the doctors stored the fertilized egg carefully so as to make sure that it wouldn't be harmed (D) or they threw it randomly around the lab in a way that made it very likely that the egg would be seriously damaged (C). Again, it's quite intuitive that my existence at a time t well after my birth (E) is evidence for D over C because $P(E|D) > P(E|C)$. But again, if we accept LP* rather than LP, then that inequality is irrelevant; what is relevant is whether $P(E|D \wedge I) > P(E|C \wedge I)$, where I presumably includes the information that I'm around at t to make observations. But this inequality is obviously false; given that I'm around at t to make observations, I'm certain to exist at t regardless of whether D or C is true, so $P(E|D \wedge I) = P(E|C \wedge I) = 1$. Again, LP* leads to the implausible result that E is no evidence for D over C .⁸

Of course, a defender of LP* could always reject these intuitions. I find it very implausible that my current existence couldn't possibly give me any evidence for or against any proposition, which is what LP* entails. Perhaps a defender of LP* finds this acceptable, but I take it to be an intuitive cost. I also take it to be a cost of LP* that it introduces an implausible epistemic distinction between the first-person stance and the third-person stance. For my existence can certainly be evidence for or against lots of propositions for *you*; after all, *you* were not guaranteed to observe my existence. But then, if

⁸ Thanks to John Roberts for helpful discussion of this sort of case.

I survive the firing squad, and you and I both see this happen, our (seemingly) shared observation of my survival is evidence for D over C for *you*, whereas it is not evidence for D over C for *me*. And while there may well be places in epistemology where some sort of relativity to epistemic agents is acceptable, I do not think that this is such a case. And so, I claim, a principle that avoids this kind of relativity would be preferable to LP*.

6.3 Probabilistic dependence between I and the hypotheses

In the firing squad case from Section 6.2, the trouble with LP* seems to be that sometimes the information I about the process of observation is *itself* evidence for one of the hypotheses under consideration, in virtue of *itself* being probabilistically non-independent of those hypotheses. In the firing squad case, the information I that I'm around at t to make observations is *itself* more likely on the supposition of D than on the supposition of C , whereas that seemingly relevant fact gets ignored when we follow LP* and simply *assume* the truth of I when evaluating the relevant likelihoods.

In cases where I is probabilistically independent of the relevant $H1$ and $H2$, as Eddington's fish case, LP* gives us the correct results. In Eddington's case, where $I =$ 'I'm using a net with large holes', I is probabilistically independent of L and S ; I'm no more or less likely to use a net with large holes on the supposition that L is true than on the supposition that S is true. In contrast, in the firing squad case, where I is itself more likely on the supposition of D than on the supposition of C , LP* runs into trouble.

It's easy to change Eddington's fish case into one where I fails to be probabilistically independent of S and L . Suppose that I have a research assistant who knows the composition of the lake (i.e. he knows whether S or L is true), and suppose that he really wants to maximize the probability that the fish I catch is from the 'majority' group (i.e. a large fish if L is true, and a small fish if S is true). So, if L is true, my assistant makes sure that I use a net with large holes, so as to maximize the chance that I catch a large fish, and if S is true, he makes sure that I use a net with small holes, to maximize the chance that I catch a small fish. And suppose that I know all of this about my research assistant's epistemic state and plan.

Now, suppose that my research assistant brings me a large-holed net and I catch a large fish. Here's a very bad argument that this provides me with no evidence for L :

My process of observation included using a large-holed net (I), so by LP* the relevant question is whether $P(E|L \wedge I) > P(E|S \wedge I)$. But here, $P(E|L \wedge I) = P(E|S \wedge I)$, since on the assumption that I'm using a large-holed net, I'm just as likely to catch a large fish regardless of whether L or S is true.

The problem with the argument above is that it ignores the probabilistic dependence between I and L . Due to the background information that I have about my research assistant's behavior, I know that the information I that I used a large-holed net is itself evidence for L , even if that information *also* introduces a selection bias which affects how we should react to E .

Now, someone might object that I haven't identified a problem here with LP* itself, since LP* only characterizes the circumstances under which E is evidence for L , and in the research assistant fish case, it's I , not E , that is the real evidence for L .⁹ Now, I'm skeptical that we can make a clean distinction between whether E or I is 'doing the real work' when the two are positively probabilistically correlated as they are in the research assistant fish case, but my real response is that I'm happy to agree that once we know I , E is no further evidence for L . But in the FTA, I don't really care whether it's E or I that constitutes evidence for D . In the FTA, I think that E is evidence for D before we take I into consideration, and then once we do take I into consideration, I 'blocks' the evidential effect of E , but only by a kind of 'preemption' that occurs in virtue of I itself being evidence for D .¹⁰ But surely, the interesting and important question about the FTA is whether we should be more confident in D after learning E through a process of observation characterized by I , not whether it's E or I that deserves the epistemic 'credit' for that increased confidence.

In the next section, I'll develop a proposal about how we should take account of information like I that will avoid the problems that we encountered with LP*, and that is designed to background this issue about whether E or I deserves the epistemic 'credit' in any particular case.

7 My Proposal

Let's go back to the research assistant fish case. If we're interested in the circumstances under which the collection of E through a process characterized by I is evidence for L , we can't use $P(E|L \wedge I)$ and $P(E|S \wedge I)$ as likelihoods, since those values will be equal, and will be insensitive to the probabilistic dependence between L and I . Again, in the research assistant fish case, it is relevant that L makes I very likely (and that S makes I very unlikely), so we might expect to somehow make use of a quantity like $P(I|L)$ and/or $P(I|S)$. Moreover, the fact that L makes I more probable than S does *increase* the extent to which L is confirmed over S , so we should expect the amount that L is confirmed over S to vary *directly* with $P(I|L)$ (or to vary *inversely* with $P(I|S)$).

⁹ Thanks to Elliott Sober for helpful discussion of this worry.

¹⁰ Which itself occurs because I and E are positively probabilistically correlated. I also think the same thing is going on in the firing squad case.

My suggestion in the research assistant fish case is that, instead of using $P(E|L \wedge I)$ and $P(E|S \wedge I)$ as likelihoods as LP* requires, we instead use $P(E|L \wedge I) \times P(I|L)$ and $P(E|S \wedge I) \times P(I|S)$. These quantities clearly take account of the probabilistic dependence between I , L , and S in the desired manner. And they clearly give us the correct result in the research assistant fish case; as already discussed, $P(E|L \wedge I) = P(E|S \wedge I)$ in this case, so $P(E|L \wedge I) \times P(I|L) > P(E|S \wedge I) \times P(I|S)$ iff $P(I|L) > P(I|S)$, which is clearly true.

More generally, my proposal is:

Likelihood Principle (LP**):** The fact that E was collected through a process of observation characterized by I is evidence for $H1$ over $H2$ iff $P(E|H1 \wedge I) \times P(I|H1) > P(E|H2 \wedge I) \times P(I|H2)$

And it's easy to see that $P(E|H1 \wedge I) \times P(I|H1) = \frac{P(E \wedge H1 \wedge I)}{P(H1 \wedge I)} \times \frac{P(I \wedge H1)}{P(H1)} = \frac{P(E \wedge H1 \wedge I)}{P(H1)} = P(E \wedge I|H1)$ (and similarly that $P(E|H2 \wedge I) \times P(I|H2) = P(E \wedge I|H2)$). So, LP** can be equivalently stated as:

Likelihood Principle (LP**):** The fact that E was collected through a process of observation characterized by I is evidence for $H1$ over $H2$ iff $P(E \wedge I|H1) > P(E \wedge I|H2)$

In other words, LP** amounts to conjoining I on the left of the conditionalization bar, rather than assuming I on the right of the conditionalization bar as LP* requires. So LP** amounts to treating I as 'just another piece of evidence', to be treated just like E , and thus doesn't introduce a fundamental distinction between information that serves as evidence and information about the process of observation. And if I is part of our evidence, as I think it is, then I have just as legitimate a claim to respecting the RTE by advocating LP** as the Objection from Anthropic Bias has to be respecting the RTE by advocating LP*.

As a referee from this journal points out to me, LP-like principles in general can be motivated as follows: we are initially interested in the posteriors $P(H1|E)$ and $P(H2|E)$. But then we notice that, if we want to isolate the evidential effect specifically of E , we have to 'factor out' potential differences in the priors $P(H1)$ and $P(H2)$. This leads us, via Bayes' Theorem, to focus on the likelihoods $P(E|H1)$ and $P(E|H2)$. But if, as I've been urging, the right way to understand the RTE is as recommending an initial interest in $P(H1|E \wedge I)$ and $P(H2|E \wedge I)$, then analogous reasoning to the above will lead us, via Bayes' Theorem, to the conclusion that the correct way to isolate the evidential effect of E is to focus on $P(E \wedge I|H1)$ and $P(E \wedge I|H2)$. And, of course, this is precisely what LP** does; in fact, I think this is another way of putting the point that LP** treats I as 'just another piece of evidence'. And the core problem

with LP* and the Objection from Anthropic Bias, it seems to me, is that it is committed to treating *I* differently to *E*.

In addition to the research assistant fish case, LP** delivers the correct verdict in all of the uncontroversial cases that we've seen.

Consider Eddington's original fish case, without the research assistant twist that I added. According to LP**, we should compare $P(E|L \wedge I) \times P(I|L)$ and $P(E|S \wedge I) \times P(I|S)$. And since *I* is independent of *L* and *S* in this case (since there's no research assistant), it's clear that $P(I|L) = P(I|S)$. And we've already agreed in this case that $P(E|L \wedge I) = P(E|S \wedge I)$, so LP** entails that we have no evidence for *L* over *S*, which is the right result. And in general, whenever *I* is probabilistically independent of the relevant hypotheses (i.e. when $P(I|H1) = P(I|H2)$), $P(E|H1 \wedge I) > P(E|H2 \wedge I)$ iff $P(E \wedge I|H1) > P(E \wedge I|H2)$, so LP** will never issue in a different result than LP*. Thus, LP** respects the observation made above that in cases where *I* is probabilistically independent of *H1* and *H2*, LP* is true. Thus, LP** explains the apparent plausibility of LP*: LP* is true when *I* is probabilistically independent of the relevant hypotheses.

In the firing squad case, LP** entails that we should compare $P(E|D \wedge I) \times P(I|D)$ and $P(E|C \wedge I) \times P(I|C)$. As already observed, $P(E|D \wedge I) = P(E|C \wedge I)$, so $P(E|D \wedge I) \times P(I|D) > P(E|C \wedge I) \times P(I|C)$ iff $P(I|D) > P(I|C)$. And since $P(I|D) > P(I|C)$ in the Firing Squad case, we have evidence for *D* over *C*, which is again the right result.

LP** also solves—or, perhaps *dissolves*—the line-drawing problem. Recall that the line-drawing problem arose because, if we include too much information about the process of observation in *I*, we'll end up with the absurd result that *E* can't be any evidence for anything. That problem resulted from a picture where we're comparing $P(E|H1 \wedge I)$ and $P(E|H2 \wedge I)$ (as LP* requires), and worrying about how to characterize *I* in a way that didn't include so much information that these quantities would always be equal (as they would be if *I* entails *E*, for then $P(E|H1 \wedge I) = P(E|H2 \wedge I) = 1$). But once we switch to comparing $P(E \wedge I|H1)$ and $P(E \wedge I|H2)$ as LP** requires, then there's no longer any concern about 'including too much information' in *I*. For suppose that *I* entails *E* by being equivalent to the conjunction of *E* with some additional information *J*. Then $P(E \wedge I|H1) = P(E \wedge E \wedge J|H1) = P(E \wedge J|H1)$ and $P(E \wedge I|H2) = P(E \wedge E \wedge J|H2) = P(E \wedge J|H2)$. So LP** results in the same likelihoods (and thus the same likelihood ratio) regardless of whether we conjoin the *E*-entailing *I* to the left-hand side or whether we just conjoin the information in *I* that 'goes beyond' *E*—i.e. *J*—to the left-hand side. As a result, there's no worry about changing the relative values of the likelihoods by 'including too much about the process of observation' in *I*; even if *I* entails *E*, conjoining *I* to the left-hand side of the likelihoods will have the same effect

as conjoining J to the left-hand side. So even if we ‘build E into I ’, no threat of triviality looms.

Finally, note that there is no implausible epistemic distinction between the first-person and third-person stances (analogous to the one introduced by LP*) that is entailed by LP**. Since my current existence is permitted by LP** to be evidence for or against propositions *for me*, there is no tension with the obvious truth that my current existence can be evidence for or against propositions *for you*.

What does LP** entail with regard to the FTA? According to the FTA, E is evidence for D over C because $P(E|D) > P(E|C)$. According to LP* and the Objection from Anthropic Bias, E fails to be evidence for D over C because $P(E|D \wedge I) = P(E|C \wedge I)$, where $I =$ ‘I’m around at t to make observations’. According to LP**, the relevant question to ask is whether $P(E \wedge I|D) > P(E \wedge I|C)$. And it’s obvious that $P(E \wedge I|D) > P(E \wedge I|C)$, since I’m much likelier to be alive and observing a life-sustaining universe at t if D is true than if C is true. Thus, according to LP**, the Objection from Anthropic Bias fails.

8 Other Views

Of course, any defender of a version of the RTE is going to claim that her view is consistent with somehow ‘taking all of the evidence into account’, and I am no different. The trouble is that different philosophers seem to disagree about what is involved in taking all of the evidence into account, and as a result disagree about what taking all of the evidence into account requires in the case of the FTA. As discussed above, Sober thinks that the correct way to take all of the evidence into account is to include all of the information about the process of observation on the right of the conditionalization bar, as required by LP*; I’ve argued that LP** is the more natural and better-motivated way to take all of the evidence into account.

Other writers, such as Neal in his ([2006]), advocate versions of the RTE; there, Neal defends ‘the principle that one should always condition on *all* evidence—not just on the fact that you are an intelligent observer, or that you are human, but on the fact that you are a human with a specific set of memories’ (Neal [2006], p. 1, emphasis in original). I of course have no objection to such a principle, and in fact LP** is fully consistent with it. But Neal’s primary interest in his ([2006]) is in figuring out when an inference to *multiple* observers is justified. On the one hand, he thinks that some cases (such as the Doomsday Problem) motivate the thought that my existence is evidence that there are multiple observers relevantly like me and, on the other, he thinks that this thought can lead to absurd results in cases such as Bostrom’s Presumptuous Philosopher (Bostrom [2002]). Neal doesn’t explicitly address

the argument from fine-tuning to design, nor does he address cases like Eddington's fish net case, where there's some prima facie plausibility to the idea that (as LP* recommends) we ought to take account of information about the process of observation differently than we take account of 'ordinary' evidence. Part of my goal in this paper has been to argue that we *can* accommodate the relevant intuitions in cases like Eddington's fish net case without embracing LP* (i.e. by accepting LP**); since LP* has implausible results that LP** doesn't have, there is reason to prefer LP**, along with its rejection of the Objection from Anthropic Bias. If Neal's version of the RTE is correct, I think that it has to be understood as an informal statement of LP**.

There are, of course, interesting issues to do with selection bias that arise when one considers hypotheses involving multiple observers. An argument that is *related* to my version of the FTA endeavors to show that *E* is evidence that there are multiple non-interacting universes, each with its own set of fundamental constants. But the issues here are complex, and I don't think it's at all obvious that the argument from fine-tuning to Design and the argument from fine-tuning to multiple universes must stand or fall together.¹¹

I conclude that the Objection from Anthropic Bias is not a successful response to the FTA. Again, this is only a very limited defense of the FTA. I haven't argued that the FTA is ultimately successful, nor have I argued that some *other* response to the FTA couldn't work. But if there is a problem with the FTA, LP* is not the key to understanding it. Rather, we should embrace LP** and treat background information as just another piece of evidence.

Acknowledgements

Thanks to Branden Fitelson, Marc Lange, Alan Nelson, Ram Neta, John Roberts, Geoff Sayre-McCord, Elliott Sober, members of John Roberts's Spring 2009 reading group on fine-tuning arguments at UNC, as well as audiences at UNC, VCU, and the 2009 Formal Epistemology Workshop for helpful discussions and questions.

*Department of Philosophy
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599, USA
kotzen@email.unc.edu*

¹¹ See (White [2000]) for an excellent discussion of some of the disanalogies between the two arguments.

References

- Bostrom, N. [2002]: *Anthropic Bias: Observation Selection Effects in Science and Philosophy*, New York: Routledge.
- Eddington, A. [1939]: *The Philosophy of Physical Science*. Cambridge University Press.
- Fitelson, B. [2011]: 'Favoring, Likelihoodism, and Bayesianism', available from: <<http://fitelson.org/flb.pdf>>.
- Leslie, J. [1989]: *Universes*, Cambridge: Routledge.
- Monton, B. [2006]: 'God, Fine-Tuning, and the Problem of Old Evidence', *British Journal for the Philosophy of Science*, **57**, pp. 405–24.
- Neal, R. [2006]: *Puzzles of Anthropic Reasoning Resolved Using Full Non-indexical Conditioning*. Technical Report No. 0607, Department of Statistics, University of Toronto; Available from: <<http://arxiv.org/pdf/math/0608592>>.
- Sober, E. [2004]: 'The Design Argument', in W. E. Mann (ed.), *Blackwell Guide to Philosophy of Religion*, Oxford: Blackwell Publishers, pp. 117–47.
- Sober, E. [2009]: 'Absence Of Evidence and Evidence of Absence: Evidential Transitivity in Connection with Fossils, Fishing, Fine-Tuning, and Firing Squads', *Philosophical Studies*, **143**, pp. 63–90.
- Weisberg, J. [2005]: 'Firing Squads and Fine-Tuning: Sober on the Design Argument', *British Journal for the Philosophy of Science*, **56**, pp. 809–21.
- White, R. [2000]: 'Fine-Tuning and Multiple Universes', *Noûs*, **34**, pp. 260–76.